# Plate localization on the Sensory Interactive Table

Joanne Spijker University of Twente P.O. Box 217, 7500AE Enschede The Netherlands j.j.spijker@student.utwente.nl

## ABSTRACT

Food intake is directly and indirectly related to many medical conditions. Getting insight in eating behaviour, especially in social settings, and being able to provide feedback to the patient in an effective manner, is still an active research area. The Sensory Interactive Table (SIT) was developed to get insight into (social) eating behaviour and to investigate how it can be used to provide helpful feedback to the user. The table can achieve this with the roughly 200 load cells and more than 8000 LEDs, which respectively can be used for activity recognition and feedback to the user. However, software development for the SIT is still in its early stages. In this research it is investigated how plates on the surface of the SIT can be localized using a simple method of averaging the locations of load cells with sufficient weight on it. First, a data set was developed by collecting data from the table itself and processing it. Then the algorithm was developed and tested on the data set, and evaluated using the metrics  $R^2$ , MAE and RMSE.

## Keywords

Sensory Interactive Table, localization, interactive dining

## 1. INTRODUCTION

Healthy eating behavior is a vital part of a healthy lifestyle. The current practice of altering eating behavior when needed is often manually tracking the intake of food, and developing a diet tailored to the person [3]. However, this disregards the often social setting of mealtimes. Social eating behavior is a real phenomenon, which can impact eating behavior both positively and negatively [6, 5, 7]. A person on a strict diet may feel like they are imposing on the person serving the food, or the subject may take another round when the group is doing that as well. This may influence the ability of a patient to stick to a diet, and it is therefore important to get more insight in how these social settings influence eating behaviour. Additionally, more research can be used to investigate how these social settings can be used to positively influence the ability to maintain a diet.

34<sup>th</sup> Twente Student Conference on IT January 29<sup>th</sup>, 2021, Enschede, The Netherlands.

Copyright 2021, University of Twente, Faculty of Electrical Engineering, Mathematics and Computer Science.



Figure 1. The Sensory Interactive Table

#### **1.1 Sensory Interactive Table**

This is where the Sensory Interactive Table (SIT) comes in. This table was designed for automatically tracking eating behavior in a social setting. The surface consists of 199 load cells and 8358 LEDs, which can respectively detect the weight distribution on the table and give feedback to the user. The load cells are located beneath hexagon shaped tiles, which are arranged in a semi circular shape. The load cell data is combined in a Teensy, which outputs it bundled over a serial port at 72Hz. For a more detailed implementation description, the reader can refer to [4].

Currently, the hardware development of the table is finished. For the software and the interaction design, there is ongoing research. For example, machine learning has been introduced to detect different eating gestures on one load cell [8], and different methods of revealing the social space of eating were explored [2]. However, these ideas cannot be implemented concretely yet, since the structure to support the detection of behavior is not yet in place. More fundamental research into the detection of objects and ultimately social structures, still needs to be performed. This research will focus on one of these fundamental open problems: the detection of plates and tracking them in real time.

This research could serve as a stepping stone for future research, as it will become easier to implement several prototypes which rely upon automatically detecting objects on the table. This way, prototypes can be implemented on the table itself instead of being explored theoretically, and this could improve the experiments performed. Ultimately, this will help progress the research in (social) eating behavior as a whole.

## 1.2 Objective and goals

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

The main research question of this research can be described as follows:

How can plates on the Sensory Interactive Table be localized and tracked in real time?

To answer this question, two solutions have been explored. One approach is implementing a simple algorithm that averages the locations of the measurements to approximate the location of plates. The other approach involves machine learning to determine the position on the table.

Since the research time is limited, the problem has been cut up into sub questions, which can be regarded as scenarios in which the designed solution can be evaluated. The research focuses primarily on the following questions:

- 1. How can one plate be localized?
- 2. How can one plate be tracked?
- 3. How can one plate that is increased in load be tracked?

In Section 2 the way this research was carried out is described. After this, in Section 3, the output of this research is described. Consecutively the implication of the results is discussed in Section 4. Finally in Section 5 conclusions will be drawn and directions for future work are described.

## 2. METHODOLOGY AND APPROACH

#### 2.1 Data collection

A data set was needed for developing and testing an algorithm that can determine the location of plates on the surface of the table. Currently, there is no comprehensive data set which can be used for developing such a system, which is why the collection of data is required.

Since the research does not focus on human activity in the sense that there has to be a human involved in the data set, the ethics committee was not notified.

For this research, the focus was on gathering data on the location of one plate on different locations on the table. To achieve this, the following steps were performed:

- 1. Numbered paper indicators were put on the table and photographed in order to determine the locations in the data preprocessing stage, see Figure 2 as an example.
- 2. A video recording was started and the number of the measurement was included to be able to match the recording to the data.
- 3. Data recording on the table was started.
- 4. The plate was put on an indicator.
- 5. About 200 grams of separate objects was added to the plate. This step was only performed in half of the measurements to diversify the data set.
- 6. The plate was removed, including the added weight if applicable.
- 7. Sensor number 0 or 198 were pressed manually in order to be able to split the data files in the data preprocessing stage.
- 8. Steps 4 to 7 were repeated for each paper indicator on the table.



Figure 2. Paper indicators on the SIT

9. The video and data recording were stopped and the measurement and remarks were written down.

The data was recorded using a Unity-based program written by the developers of the SIT. This program does not output the raw data, but rather takes a short calibration period and outputs the calculated weight in grams. For the purpose of this research, this is a good starting place as it removes the necessary calibration step needed anyway. However, as this data is also smoothed in a certain way, machine learning application may benefit from using the raw data, as will be discussed in Section 4.

#### 2.2 Data preprocessing

In order to use the data in any application, it needed to be preprocessed and formatted. The first step was to extract the actual locations of the plates from the photos taken of the indicators. It was decided to overlay a virtual 450 by 450 unit grid over the table. This way, the coordinates of the middle points of the hexagonal tiles could be saved to integers, while still keeping the granularity low enough for our applications. To get an idea of the scale of this grid, the height of a tile is 30 units and the width is 26.

The next step was to split the recorded data files into individual files of one action each. This was done by splitting the files based on the spikes sensors 0 or 198 outputted, caused by pressing on this tile each time an individual measurement was finished. In this step, the locations were added to the individual measurement files.

After this, the data was filtered using a low pass filter. This filter is able to filter out high frequency data and let low frequency data through. One example of a low pass filter is to simply average value n with value n-1, or more previously measured values if desired. The more values are averaged, the more smoothed out the data will be. This has the drawback that it will take longer for new values to be propagated, and that it will therefore take longer to make an accurate prediction about the location of the plates. However, this predicted location is less likely to be influenced by noise in the data.

## 2.3 Analysing raw data

To get an idea what the data looks like without needing to plot all the graphs of the measurement, a simple data visualiser was built. As input it takes the raw data extracted from the table, which it then can play back. It uses a grid of circles, each circle representing a hexagon tile with a sensor, giving it a color depending on the amount of pressure it holds in this moment. This enables the observation Weight map SIT – C ×

Figure 3. A screenshot of the data visualiser

of the data in 'real time' and gives an idea of the spatial layout. From these visualisations it was clear that the algorithm proposed in the next section could be feasible in practice.

As is visible in Figure 3, it is very clear where the plate is located at that point in time. Next to this, when running the visualiser, it becomes very clear that the sensor on the bottom left is defective as it has a continuous high output.

#### 2.4 Developing an algorithm

The developed algorithm performs the following steps for each measurement:

- 1. *Tile extraction* In this step all the tiles on which there is sufficient pressure are extracted, which is accurate enough for this application since plates have quite some weight. The output is a list of tiles that have an observed pressure higher than a determined threshold. The threshold was chosen so that it performs well on this data, but this might need to be calibrated in a different setting.
- 2. Predict whether there is a plate When the amount of tiles outputted by step 1 is two or more, the algorithm will continue with step 3. If not, the algorithm will output that there is no plate on the table at the moment.
- 3. Location estimate The location is then estimated based on the weighed average location of the tiles for both the x and y axis. For example, if a tile at location (0,0) and a tile at location (10,10) have a respective weight of 1 and 2, it will estimate that the location is (6.6, 6.6). Of course, this becomes more accurate as more tiles are added to the calculation.
- 4. Saving results The predicted values are stored together with the actual values so that the evaluation metrics can be calculated easily.

#### 2.5 Testing algorithm accuracy

To test this algorithm, the estimates were tested against the actual values. Common metrics to evaluate these prediction were used:  $R^2$ , Mean Absolute Error (MAE) and



Figure 4. Example of data outputted by the SIT for one measurement without added weight

Root Mean Square Error (RMSE), as is described well by Towards Data Science in article [1].

 $R^2$  assigns a percentage to how well the model fits the data. The lower the value, the less likely it is that the model predicts a value that is close to the ground truth. It is calculated as follows:

$$R^{2} = 1 - \frac{SS_{Regression}}{SS_{Total}} = 1 - \frac{\sum(y_{i} - \hat{y}_{i})^{2}}{\sum(y_{i} - \bar{y}_{i})^{2}}$$

The Mean Absolute Error is the average of the error in prediction. It is the most clear indicator of how close the predictions on average are to the ground truth. MAE treats all errors the same, which is why large errors can influence the average significantly and sketch a skewed look on the majority of the predictions. In general, the lower the values, the better. It is calculated as follows:

$$MAE = \frac{1}{N} \sum_{i=1}^{N} |y_i - \hat{y}_i|$$

The Root Mean Square Error, in contrast to  $\mathbb{R}^2$ , gives an absolute measure instead of a relative measure of how well the model fits. Since it uses the square of the error, it is not usable to draw direct conclusions about the goodness of fit the way MAE can be used. Instead, it can be used to compare different cases. Compared to MAE, this metric gives a higher penalty to larger errors. It was chosen to use RMSE instead of just the Mean Squared Error, since this value is smaller and therefore easier to interpret. The lower the value, the better. It is calculated as follows:

$$RMSE = \sqrt{\frac{1}{N} \sum_{i=1}^{N} (y_i - \hat{y}_i)^2}$$

From these metrics, conclusions can be drawn about the effectiveness of the algorithm, and about the effectiveness of the data preprocessing.

## 3. RESULTS

#### 3.1 Data set

The collected data consists of 232 individual cases, of which half do not have added weight. One example of the raw data of such a case is visible in Figure 4. In this figure, as well as in Figure 5 and 6, the time is on the x-axis, and



Figure 5. Example of a low pass filter applied to the data in Figure 4



Figure 6. Example of data outputted by the SIT for one measurement with added weight

the output of each sensor is listed on the y-axis. This example is 90 measurements long, which in reality is about 4 seconds.

Each colored line is an individual sensor, and it is clear that most sensors have a steady output value between -1 and 1. The difference between the measurements without and with added weight is visible when comparing Figure 4 to 6. The figures also indicate that two sensors are defective. The sensors corresponding to the orange line and the blue line on respectively 11 and -10 have therefore been filtered out when processing the data.

#### 3.2 Data preprocessing

As is visible in the plots of the data output, the data contains spikes where in reality nothing is happening on the table. In order to filter this out, some data processing was needed, which was done in the form of a low pass filter. A variant that averages the 10 last observed applied to the data in Figure 4 is visible in Figure 5.

#### **3.3** Algorithm performance

The findings of applying a low pass filter to the data set and calculating the performance of the predictions on the x- and y-axis using the metrics described in Section 2.5 are listed in Table 1 and 2 respectively.

## 4. **DISCUSSION**

#### 4.1 Data set

There are several factors that make this data set not perfect for this application and that could be improved in future data collection efforts. First of all, the size of the data set is limited. Even though the method was designed to cover the entire table area, the number of data points is not enough for machine learning applications in a production setting.

The next drawback of this data set is that it currently only contains data on the location of one plate at a time. Due to COVID-19 circumstances it was not feasible to collect more data with multiple objects on the table. Data with multiple objects could help develop methods to discern between objects and tracking those as they increase and decrease in weight. It is plausible that the methods developed in this research will also work with different objects such as cups and pots since there has been no focus on detecting only plates. The weight that was added in half of the measurements therefore also did not contribute much to the research. However, this could be useful in future applications such as detecting how much food a plate carries.

Lastly, it is visible in the data that two sensors malfunction. While this can be manually avoided this time, it implicates that measurements in the area of those sensors (which are close by each other) are not reliable.

While these drawbacks are important to take into account for future data collection efforts, the data is still very usable for investigating how the load cells behave. As visible in the plots in Figures 4 and 6 the data is quite clear in these clear-cut situations.

#### 4.2 Algorithm performance

The developed algorithm in its current form is very simple and with the right analysis it can most likely be improved in several ways. However, given the simplicity of the algorithm it is promising that the Mean Absolute Error is only about 17 for the x-axis and about 9 for the y-axis with a low pass filter that averages the last 10 values.

As each tile only has a width of 26 and a height of 30, the error is on average smaller than both the width and the height of a tile. Especially considering the plates used generally cover more than 5 tiles, this is a prediction precise enough to be usable.

It is clear that the predictor has different performances for the x- and y-axis, which is probably explainable by the way the tiles are arranged. The hexagon tiles can either be arranged in a grid with a side at the top or with a corner pointing up. In this application the tiles are arranged with the corner pointing to the top. While this difference seems like a logical explanation for the difference in performance, the data has not been analysed on this issue.

The performance is not yet perfect, and this is also partly explainable by characteristics of the hardware itself. For example, the surface of the table is not perfectly flat, i.e. that some tiles are a little higher than other tiles. When a plate is put on multiple sensors of which some are higher than others, the lower tiles may not be triggered at all, causing unbalanced data points. The developers of the SIT expect that this problem will persist, also in future versions of the device. Therefore, using Machine Learning may offer a solution, since that way the classifier may learn the specific characteristics of the table, which are less practical to implement in a custom algorithm.

Lastly, the results indicate that taking the average of more measurements does not necessarily increase performance. In fact, the performance seems to drop slightly when the average of the last eight or more measurements is taken. This indicates that in applications that run on the hard-

Table 1. Results of prediction on the x-axis after applying different low pass filters to the data set

Dataset	$\mathbf{R}^2$	Mean Absolute Error	Root Mean Squared Error
Original	0.896	18.236	31.536
Low pass filter last 2	0.909	17.531	29.497
Low pass filter last 4	0.921	17.032	27.542
Low pass filter last 6	0.926	16.942	26.690
Low pass filter last 8	0.919	17.449	27.773
Low pass filter last 10	0.916	17.712	28.361
Low pass filter last 15	0.911	18.060	29.129
Low pass filter last 20	0.906	18.369	30.008
Low pass filter last 25	0.904	18.484	30.209
Low pass filter last 30	0.900	18.851	30.894

Table 2. Results of prediction on the y-axis after applying different low pass filters to the data set

Dataset	$\mathbf{R}^2$	Mean Absolute Error	Root Mean Squared Error
Original	0.922	12.604	26.314
Low pass filter last 2	0.951	11.053	21.000
Low pass filter last 4	0.976	9.618	14.589
Low pass filter last 6	0.979	9.386	13.820
Low pass filter last 8	0.977	9.517	14.358
Low pass filter last 10	0.978	9.530	13.995
Low pass filter last 15	0.977	9.692	14.268
Low pass filter last 20	0.977	9.720	14.271
Low pass filter last 25	0.977	9.687	14.214
Low pass filter last 30	0.977	9.738	14.270

ware of the SIT, fewer values need to be stored for the purpose of applying a low pass filter, which could improve performance if needed.

#### 5. CONCLUSION AND FUTURE WORK

### 5.1 Conclusion

The research performed offers an insight in the way data output can be used for localization of plates on the SIT. To be able to answer all the research questions posed in Section 1.2, more research is needed. For the first research question an algorithm was developed and evaluated, and the results are good enough for real life applications. However, as this research was focused on detecting one object on the table, this is not ready to be applied in a wellrounded application. Regardless, important characteristics of the SIT were discovered and the data collected and findings from this research can be used as a stepping stone for future research.

## 5.2 Future Work

Early tests in this research have indicated that localization of objects on the SIT may benefit greatly from using machine learning on a well developed data set. Since the input for such a model will contain the same characteristics as the real life input, the model will be able to adjust to these inconsistencies more adequately than a theoretical algorithm can. During the data analysis phase, experiments using Machine Learning have been performed with promising results, but this requires more research as the experiments that were performed were not carried out in a scientific manner. In order to train a model that is capable of accurately predicting such locations, a larger and more diverse data set is needed.

#### 6. ACKNOWLEDGEMENTS

I would like to thank my supervisor dr.ir. Haarman and track chair dr. Theune for their support and guidance during the research. Their remarks and kind words have helped me tremendously, in particular during these difficult COVID-times. Next to this, I would like to thank Justin Adriani and Sieta de Jong for being great sparring partners and providing moral support, and Justin especially for helping me out with taking measurements.

## 7. REFERENCES

- 3 best metrics to evaluate regression models. https://towardsdatascience.com/what-are-the-bestmetrics-to-evaluate-your-regression-model-418ca481755b. Accessed: 2021-01-24.
- [2] R. de Vries, J. Haarman, E. Harmsen, D. Heylen, and H. Hermens. The sensory interactive table: Exploring the social space of eating. 11 2020.
- [3] A. B. J. M. Guendalina Graffigna, Serena Barello. The motivating function of healthcare professional in ehealth and mhealth interventions for type 2 diabetes patients and the mediating role of patient engagement. 2016.
- [4] J. Haarman, R. de Vries, E. Harmsen, H. Hermens, and D. Heylen. Sensory interactive table (sit): Development of a measurement instrument to support healthy eating in a social dining setting. *Sensors (Switserland)*, 20(9), May 2020. This article belongs to the Special Issue Sensor-Based Activity Recognition and Interaction.
- [5] S. Higgs. Social norms and their influence on eating behaviours. Appetite, 86:38–44, 2015.
- [6] S. Higgs and H. Ruddock. Social Influences on Eating, pages 277–291. Springer International Publishing, Cham, 2020.
- [7] D. A. Roth, C. P. Herman, J. Polivy, and P. Pliner. Self-presentational conflict in social eating situations: a normative perspective. *Appetite*, 36(2):165–171, 2001.
- [8] D. A. Uğurlu. The feasibility of classification of eating-related load cell signals. August 2020.