KU LEUVEN

ARENBERG DOCTORAL SCHOOL
Faculty of Engineering Technology

Automatic Food Intake Monitoring For The Ageing Population

Gert Mertes

Supervisors:
Prof. dr. ir. B. Vanrumste
Prof. dr. ir. H. Hallez
Prof. dr. ir. ing. T. Croonenborghs

Dissertation presented in partial fulfilment of the requirements for the degree of Doctor of Engineering Technology (PhD)

January 2019

Automatic Food Intake Monitoring For The Ageing Population

Gert MERTES

Examination committee:
Prof. dr. ir. B. Lievens, chair
Prof. dr. ir. B. Vanrumste, supervisor
Prof. dr. ir. H. Hallez, supervisor
Prof. dr. ir. ing. T. Croonenborghs, supervisor
Prof. dr. ir. T. Goedemé
Prof. dr. C. Matthys
Prof. dr. ir. W. Chen
(Fudan University, China)
Prof. dr. D. Beeckman
(Ghent University)
Dr. L. Cuypers
(COMmeto)

Dissertation presented in partial fulfilment of the requirements for the degree of Doctor of Engineering Technology (PhD)

January 2019

© 2019 KU Leuven – Faculty of Engineering Technology
Uitgegeven in eigen beheer, Gert Mertes, Andreas Vesaliusstraat 13, B-3000 Leuven (Belgium)

Alle rechten voorbehouden. Niets uit deze uitgave mag worden vermenigvuldigd en/of openbaar gemaakt worden door middel van druk, fotokopie, microfilm, elektronisch of op welke andere wijze ook zonder voorafgaande schriftelijke toestemming van de uitgever.

All rights reserved. No part of the publication may be reproduced in any form by print, photoprint, microfilm, electronic or any other means without written permission from the publisher.

Voorwoord

Zo, hier zijn we dan. De culminatie van 4 jaar onderzoek met hier en daar een traan en liters zweet, maar gelukkig geen bloed. Dat is wat er voor u ligt in de vorm van een bergje papier. Of misschien leest u digitaal en is al dat gereduceerd tot een wirwar van eentjes en nullen - slechts enkele megabytes. Dat klinkt haast poëtisch. Poëzie was echter nooit mijn sterkste vak. Dingen uit elkaar halen, dáár was ik van jongs af aan goed in (de familie zal dat, tot spijt, kunnen beamen). Ergens is dat ook de essentie van dit werk. De concepten, ideeën en algoritmes uit deze thesis zijn tientallen - zo niet honderden - keren opgebouwd, om ze vervolgens terug uit elkaar te halen en opnieuw te beginnen. Ik kijk met trots terug op het verrichte onderzoek dat geleid heeft tot deze tekst. Het doel van een doctoraat is om gedeeltelijk zelfstandig te werken, maar deze thesis had er niet geweest zonder cruciale bijdragen of hulp van meerdere personen. In dit voorwoord wil ik dan ook deze personen bedanken.

Allereerst een welgemeend bedankje aan mijn promotoren Bart Vanrumste, Hans Hallez en Tom Croonenborghs. Bart, bedankt dat je me deze kans gaf. Toen ik de eerste keer bij jou op gesprek kwam was ik enorm onzeker om in te tekenen voor een doctoraatstraject. "Vier tot vijf jaar, da's toch lang hé", dacht ik toen. Van die onzekerheid blijft vandaag echter geen druppel over en ik kijk met plezier terug op tijd die ik onder jouw begeleiding heb doorgebracht. Oprecht bedankt voor de vrijgemaakte tijd, nalezen van papers en onze talrijke technische discussies. Hans, mede dankzij jou kon ik financiering verkrijgen om gedurende mijn hele doctoraat ongestoord aan mijn onderzoek te werken. Hiervoor ben ik je zeer dankbaar. Verder heb ik genoten van onze samenwerking en buitenlandse verplaatsingen die we samen hebben gemaakt. Tom, ik kan met zekerheid zeggen dat jouw feedback de kwaliteit van mijn papers naar een hoger niveau heeft getild. Zelfs op de valreep van een deadline kon ik rekenen op jouw inzicht en gaf je nog aan hoe ik op korte tijd de paper toch nog kon verbeteren. Hiervoor en voor onze fijne samenwerking wil ik je dan ook oprecht bedanken.

Verder wil ik de examencommissie bedanken om te zetelen als jurylid van dit doctoraat. Dr. Ludo Cuypers, bedankt om sinds de start van het doctoraat betrokken te zijn en hartelijk dank voor de fijne samenwerking. Prof. dr. ir. Toon Goedemé verdient een extra bedankje want dankzij jou kwam ik als prille ingenieur in contact met Bart, en de rest is geschiedenis. Verder bedankt aan Prof. dr. ir. Bart Lievens, Prof. dr. Christophe Matthys, Prof. dr. ir. Wei Chen en Prof. dr. Dimitri Beeckman voor de vrijgemaakte tijd.

Furthermore, I would like to extend additional gratitude to Wei Chen for receiving me for a research visit at Fudan University in Shanghai. These six months are one of my most cherished and valuable experiences as a PhD student. Thank you for taking the time out of your busy schedule to work with me on my research in China. This research would furthermore not have been possible without Prof. dr. Jia Jie and dr. Li Ding. Thank you for receiving me at Huashan hospital and working with me during the data collection process. Parts of this thesis were made possible thanks to your collaboration, for which I am extremely grateful. Furthermore, I would like to thank the doctors, students, nurses and patients of Jing'An Huashan hospital in Shanghai for their support during the research. A final thanks goes to the staff and students of the Center for Intelligent Medical Electronics at Fudan University for their hospitality and pleasant collaboration during my stay in Shanghai.

I'd also like to thank Alina for being my go-to spell checker.

Tijdens mijn doctoraat had ik het genoegen om deel uit te maken van enkele projectteams. Allereerst een bedankje aan de collega's die deelnamen aan het FallRisk project: Femke De Backere, Jelle Nelis, Shirley Elprama en Jan van den Bergh: dank je voor de toffe samenwerking. Hoewel het onderzoek verricht tijdens dit project niet rechtstreeks in deze tekst is verwerkt, heb ik er toch veel uit geleerd en hiervoor ben ik het hele projectteam dankbaar.

Het project dat ongetwijfeld aan de fundering ligt van heel dit doctoraat is het Ingenieurs@WZC project. De collega's die hebben meegewerkt aan dit project ben ik dan ook zeer dankbaar. Bedankt aan Peter Bernaers om de deuren van het WZC open te zetten en Jessica Hekking en Patricia Sabbe voor de ondersteuning tijdens het project. In het bijzonder dank aan Marijs Meul, An Mondelaers en Kristien Rombouts van het *projectteam food intake*. Ik heb enorm genoten van onze samenwerking en de co-creatie sessies die we samen hebben georganiseerd. Wat een top team! Grote delen van dit doctoraat zijn verricht dankzij jullie hulp en hiervoor ben ik jullie immens dankbaar. Verder ook mijn oprechte bedankjes aan het personeel en de inwoners van WZC Edouard Remy. Dankzij jullie bereidheid in het testen van de technologie en verzamelen van data is dit werk tot stand kunnen komen, dus een bedankje kon hier uiteraard niet ontbreken.

De vele collega's met wie ik het genoegen had om samen te werken verdienen natuurlijk ook een bedankje. Allereerst een bedankje voor de toffe samenwerking aan de collega's van KUL Campus Geel, waaronder Patrick Colleman, Vic Van Rooie en Staf Vermeulen. Jeffrey en Bram, bedankt voor de aangename samenwerking, het scheppen van sfeer op *onze* bureau en de extracurriculaire activiteiten tijdens de middagpauze. Verder bedankt aan de ex-collega's van ADVISE: Peter, Stijn, Gert D., Lode en Bert. Ook de huidige collega's van e-Media Lab verdienen een bedankje in het bijzonder, waaronder, maar zeker niet beperkt tot, Luc, Vero, Ine, Hannelore, Jonas, Karsten, Dimitri, Benjamin, Yiyuan en Duowei en alle andere doctoraatsstudenten en medewerkers van KUL Campus Groep T. Mede-doctorandus Greet Baldewijns verdient een bedankje in het bijzonder. Greet, ik herinner me nog als gisteren de dag dat je me kwam afhalen aan het onthaal op mijn eerste werkdag en de begeleiding die ik van je heb gekregen om me wegwijs te maken. Ik heb je zien groeien van doctoraatsstudent tot doctor, we hebben afgezien op conferenties die veel te lang duurden en zijn letterlijk gevallen en opgestaan. Hartelijk dank voor de fijne samenwerking.

Dit alles had ook niet mogelijk geweest zonder de immense steun van mijn mama en papa: Viviane en Ivo. Dankzij jullie aanmoediging kon ik mijn studies en doctoraat succesvol afronden. Jullie stonden elke dag - zelfs vandaag nog - klaar met een bordje fruit of aanmoedigend woordje. Ook voor het nastreven van mijn hobby's kon ik steeds op jullie ondersteuning en aanmoediging rekenen. Hiervoor ben ik jullie oneindig dankbaar. Verder nog een bedankje aan mijn zus Kirsten om er te zijn wanneer ik het nodig had.

Last, but certainly not least, I would like to thank my partner, Jing, for supporting me with patience and love.

Abstract

Food intake monitoring can play an important role in the prevention of malnutrition among older adults. Traditional monitoring methods typically involve the use of pen-and-paper food diaries or questionnaires. While digital alternatives exist, these tools rely on manual data entry, often multiple times a day. Furthermore, the recorded data may be incomplete and contain mistakes due to human error or a deliberate misreporting of the food intake. While these methods are considered to be the gold standard for food intake, they are rarely used in the ageing population due to their time consuming nature. Nevertheless, nutrition plays an important role in the health status of older adults. Malnutrition has been linked with decreased muscle strength, poorly healing wounds, increased hospital admission time and increased hospital mortality rate. Preventing malnutrition by means of a targeted nutritional intervention can prevent these problems and increase general quality of life. A routine monitoring for early recognition and treatment of malnutrition should therefore be included in the routine care of older adults.

Technology can play an important role in the food monitoring process. Sensors may be employed to automatically measure eating activity or the amount of consumed food, which can supplement traditional methods, lowering the recording burden on the user and caregiver. Several sensor systems have been proposed in the literature to accomplish this task, from wearable devices to table embedded scales and camera based methods running on a smartphone. Research specifically investigating methodologies for use in the older population, however, remains sparse. The wearable systems proposed in the literature may not be comfortable to use for longer periods of time or be stigmatising to the older adult. Table embedded scales require an extensive adaptation of the existing eating surface, while smartphone usage has proven difficult in older adults. There is an acute need for tools that are comfortable to use by older adults to aid in the daily care and prevention or treatment of malnutrition or related disease.

Two sensor systems and corresponding algorithms are proposed in this thesis. The first is an accelerometer based wearable system, with the accelerometer mounted on the eyeglasses of the user. The eyeglasses are used in this context as a platform to mount the accelerometer in close proximity to the head. The ears of the glasses are able to transmit the vibrations and movements of the mastication muscles during chewing to the accelerometer, where it can be converted into an acceleration signal and used for the detection of chewing activity. A chewing detection algorithm is proposed based on supervised machine learning techniques. Data was recorded from older adults in a nursing home to train and validate the model. The results show that an accelerometer worn this way can be used to detect chewing activity.

Furthermore, a smart plate system is presented. For the purpose of the research, a prototype was designed and developed consisting of a custom embedded system and sensors. The system consists of a base station and an off-the-shelf polymer plate that is mounted on top of the base. Weight sensors in the base accurately measure the weight of consumed food from the plate. The novelty of this system is the ability to measure the location of individual bites on the plate. In combination with a compartmentalised plate, the system can estimate from which compartment a bite was taken, without any sensors or electronics embedded into the plate itself or physical changes to the plate. All hardware is located into the base station and is separate from the plate. For the bite localisation to work, an accurate detection of the individual bites is required. For this, a novel bite detection algorithm is presented based on a supervised learner. Data was recorded from older adults eating a meal with the plate in a nursing home and hospital, and is used to train and validate the model. Results show that the system works as expected, with a bite detection algorithm that improves on the state of the art and the ability to measure the consumed food per compartment. With prior knowledge of which food type was served in which compartment, this can allow for an automatic estimation of the total amount of ingested calories.

Finally, an exploratory study into behavioural analysis using the smart plate is presented. We show that parameters extracted from individual bites detected with the smart plate may be used as a descriptor of behavioural traits during eating.

The systems and algorithms presented in the thesis have the potential to lower the threshold for the adaptation of sensor based food intake monitoring in older adults. Furthermore, the tools can be employed for research in other target groups, from prevention and treatment of obesity to quantified-self applications in young to middle aged adults.

Samenvatting

Het opvolgen van het voedingspatroon is belangrijk in het voorkomen van ondervoeding bij oudere personen. Traditionele manieren om het voedingspatroon op te meten, ook wel voedselmonitoring genoemd, werken typisch met pen-en-papier voedseldagboeken of vragenlijsten. In een voedseldagboek wordt de hoeveelheid voedsel dat dagelijks geconsumeerd wordt nauwkeurig bijgehouden. Hoewel digitale versies hiervan bestaan, moeten ze manueel ingevuld worden door de persoon of een mantelzorger, vaak meerdere keren per dag. De ingevulde informatie kan onvolledig zijn of fouten bevatten. Deze methodes gelden als gouden standaard voor het opvolgen van het voedingspatroon, maar ze worden door hun arbeidsintensiviteit zelden gebruikt door ouderen. Zorgpersoneel heeft typisch ook niet de tijd om hierbij te helpen. Dit terwijl voeding een belangrijke parameter is in de gezondheidsstatus van ouderen. Ondervoeding is gelinkt aan verminderde spiersterkte, slecht genezende wonden en een verhoogde ziekenhuisopnametijd en sterftecijfer. Door ondervoeding te vermijden kunnen deze problemen verminderen en kan de algemene levenskwaliteit van de oudere worden verbeterd. Een routine monitoring voor de herkenning van ondervoeding zou opgenomen moeten worden in de dagelijkse zorg van oudere personen.

Technologie kan hierin een belangrijke rol spelen. Sensoren kunnen worden gebruikt om eetactiviteit en hoeveelheid automatisch op te meten. Deze informatie kan gebruikt worden om het traditioneel voedseldagboek automatisch in te vullen, wat de werklast voor de gebruiker vermindert. Verschillende soorten sensoren zijn reeds voorgesteld in de literatuur om deze taak uit te voeren, van draagbare activiteit trackers, tot een weegschaal ingebouwd in de eettafel of camera gebaseerde methodes in de vorm van een app op de smartphone. Onderzoek dat toegespitst is specifiek op oudere personen is echter beperkt. De draagbare systemen uit de literatuur zijn niet altijd comfortabel voor langdurig gebruik en kunnen stigmatiserend zijn. Een weegplatform ingebouwd in de tafel vergt de nodige aanpassingen aan de infrastructuur en ouderen kunnen niet altijd even goed overweg met een smartphone. Er is een nood aan tools die

makkelijker en comfortabeler zijn voor gebruik door oudere personen voor het opmeten van het dagelijks eetpatroon.

In deze thesis worden twee systemen en bijhorende algoritmes voorgesteld. Het eerste is een draagbaar systeem met een accelerometer bevestigd op de bril van de persoon. De bril dient om de sensor te bevestigen in de nabijheid van het hoofd. Het montuur van de bril maakt rond de oren contact met de spieren die bewegen tijdens de kauwbeweging. De trillingen en bewegingen van het kauwen worden via het montuur doorgegeven aan de accelerometer, waar het wordt omgezet in een signaal waarmee kauwbeweging is te onderscheiden van andere activiteiten. Op deze manier kan eetactiviteit automatisch gemeten worden. Het algoritme is gebouwd en getest met data opgenomen bij ouderen in een woon-zorgcentrum. De resultaten tonen aan dat kauwbeweging op deze manier gedetecteerd kan worden.

Verder wordt er een slim bord getoond. Een prototype van dit slim bord werd gebouwd tijdens het onderzoek. Het systeem omvat een kunststof bord met 3 vakken dat gemonteerd wordt op een basisstation. Sensoren in het basisstation meten het gewicht van het eten op het bord tijdens de maaltijd. Hiermee is de hoeveelheid opgegeten voedsel te meten. Het vernieuwende aspect van dit systeem is dat individuele happen kunnen gedetecteerd en gelokaliseerd worden op het bord. Happen worden toegewezen aan het vak waaruit ze genomen zijn. Indien men op voorhand weet welk type voedsel er in elk vak ligt, kan het aantal opgenomen calorieën worden bepaald per vak en voor de totale maaltijd. Alle elektronica en sensoren zitten verwerkt in het basisstation, het kunststof bord wordt enkel gemonteerd op de sensoren zonder verdere aanpassing. Data van ouderen in een woon-zorgcentrum en ziekenhuis werd gebruikt voor het bouwen en testen van een algoritme om individuele happen te detecteren. De resultaten tonen aan dat het systeem werkt zoals verwacht en het algoritme is in staat om happen te detecteren en lokaliseren tijdens normale maaltijden.

Tot slot wordt er een studie voorgesteld die onderzoekt of het mogelijk is om extra parameters te extraheren uit de happen gedetecteerd met het slim bord. Deze parameters kunnen worden gebruikt om bepaalde gedragspatronen te herkennen tijdens de maaltijd, wat interessant kan zijn voor studies naar het gedrag van personen tijdens het eten.

Het werk dat voorgesteld wordt in deze thesis heeft het potentieel om het toepassen van voedselmonitoring bij ouderen te vergemakkelijken. Verder is er een mogelijkheid tot het gebruiken van de technologie in een breder toepassingsgebied bij jongere gebruikers.

Contents

Abstract v

Samenvatting vii

Contents ix

List of Abbreviations xvi

List of Figures xvii

List of Tables xxi

1 Introduction 1

1.1 Malnutrition in the Ageing Population 1

1.1.1 Self-reporting 1

1.1.2 Nutritional Screening Tools 2

1.1.3 Automatic Food Intake Monitoring 3

1.2 Research Goals and Motivation 3

1.3 Outline of the Thesis 4

2 Background 7

2.1 Introduction 7

2.2 Related Works 8
2.2.1 Wearable Systems 8
2.2.2 Imaging Systems 9
2.2.3 Table Embedded Scale 11
2.3 Machine Learning Techniques 13
2.3.1 Data Representation 13
2.3.2 Validation 13
2.3.3 Support Vector Machine 14
2.3.4 k-Nearest Neighbours 16
2.3.5 Decision Tree Learning 17
2.3.6 Feature Selection 18
2.3.7 Evaluation Metrics 20
2.4 Conclusion 22

3 Food Intake Monitoring for Older Adults: Stakeholders, Needs & Requirements 23

3.1 Introduction 23
3.2 The Engineers@Care Homes project 23
3.3 Stakeholder Study: Focus Groups 25
3.3.1 Older Adults 25
3.3.2 Healthcare professionals 27
3.4 Ageing Population in China 29
3.5 Outline and Rationale of the Proposed Solution 30

4 Data Acquisition 33

4.1 Introduction 33
4.2 Capture Setup & Protocol 33
4.2.1 Wearable System 33

4.2.2 Smart Plate . . . 34
4.2.3 Measurement Protocol . . . 35
4.2.4 Labelling . . . 36
4.2.5 Ethics . . . 37
4.3 Datasets . . . 38
4.3.1 Belgium . . . 38
4.3.2 China . . . 40
4.4 Conclusion . . . 41

5 Detection of Chewing Using a Glasses Mounted Accelerometer 45
5.1 Introduction . . . 46
5.1.1 Anatomy of Chewing . . . 47
5.2 Methods . . . 47
5.2.1 Accelerometer . . . 48
5.2.2 Data . . . 50
5.2.3 Pre-processing . . . 53
5.2.4 Feature Extraction . . . 54
5.2.5 Feature Selection . . . 55
5.2.6 Training and Validation . . . 57
5.3 Results . . . 58
5.4 Discussion . . . 60
5.5 Strengths and Limitations . . . 62
5.6 Chewing Strength . . . 63
5.7 Conclusion . . . 65

6 Design of a Smart Plate 67
6.1 Introduction . . . 67
6.2 Bite Localisation . . . 69

6.3 Strain Gauge Load Cell . . . 72
6.4 Mechanical Design . . . 74
6.5 Electrical Design . . . 75
6.5.1 Embedded System . . . 76
6.5.2 Power Supply . . . 77
6.5.3 Load Cell Preamp . . . 77
6.5.4 PCB . . . 78
6.6 System Tests . . . 79
6.6.1 Calibration . . . 79
6.6.2 Linearity . . . 79
6.6.3 Drift . . . 79
6.6.4 Repeatability . . . 80
6.6.5 Position Accuracy . . . 80
6.7 Conclusion . . . 80

7 Detection of Food Bites Using a Smart Plate During Unrestricted Eating 83

7.1 Introduction . . . 84
7.2 Preliminary Study . . . 85
7.3 Challenges During Unrestricted Eating . . . 87
7.3.1 Bite Types . . . 89
7.4 Methods . . . 91
7.4.1 Data . . . 91
7.4.2 Stability Detection . . . 92
7.4.3 Bite Classifier . . . 92
7.4.4 Detection Pipeline . . . 95
7.4.5 Performance Metrics . . . 95
7.5 Results . . . 97

7.6 Discussion 100
7.7 Strengths and Limitations 102
7.8 Conclusion 103

8 Quantifying Eating Behaviour With a Smart Plate 105
8.1 Introduction 105
8.2 Bite Parameters 106
8.3 Age Specific Behaviour 108
8.4 Person Specific Behaviour 110
8.4.1 Data & Methods 110
8.4.2 Results & Discussion 112
8.5 Conclusion 114

9 Conclusion & Future Work 115
9.1 Summary of Results 115
9.2 Future Work 117
9.2.1 Accelerometer Based Monitoring 118
9.2.2 Smart Plate Based Monitoring 119
9.2.3 Behavioural Analysis 120
9.2.4 Valorisation Potential 120

Bibliography 123

List of Publications 135

List of Abbreviations

ADC	Analogue to Digital Converter
ADL	Activities of Daily Living
AVDD	Analogue Voltage Supply
CIRS	Cumulative Illness Rating Scale
CNN	Convolutional Neural Network
COP	Centre of Pressure
DC	Direct Current
EMG	Electromyography
FFT	Fast Fourier Transform
FIR	Finite Impulse Response
FMA	Fugl-Meyer Assessment
FN	False Negative
FP	False Positive
GNRI	Geriatric Nutritional Risk Index
GP	General Practitioner
IMU	Inertial Measurement Unit
IQR	Interquartile Range
kNN	k-Nearest Neighbours
LOPO	Leave-One-Person-Out

LPF Low Pass Filter

MAV Mean Absolute Value

MCU Microcontroller Unit

MEMS Micro-Electro-Mechanical Systems

MIB Monitoring of Ingestive Behaviour

MNA Mini Nutritional Assessment

MRMR Maximum Relevance Minimal Redundancy

PCB Printed Circuit Board

PPV Positive Predictive Value

RF Random Forest

RFID Radio-Frequency Identification

RTC Real Time Clock

SPP Serial Port Profile

SVM Support Vector Machine

TN True Negative

TP True Positive

UEM Universal Eating Monitor

USB Universal Serial Bus

List of Figures

1.1 Schematic overview of the thesis structure. 4

2.1 The Automatic Ingestion Monitor with jaw strain gauge, accelerometer module, wrist mounted hand-to-mouth proximity sensor and smartphone. 9

2.2 Image segmentation and volume estimation of different food types. 10

2.3 Implementation of a table embedded scale. Food is served on a tray that is in turn placed on top of the weighing scale. The weight data from the scale can be used to detect individual bites during the meal. 12

3.1 Focus group with older adults about activity trackers and food intake monitoring technologies during the *Engineers@Care Homes* project. 27

3.2 Mock-up of what a personalised food intake report could look like when presented to healthcare professionals. 28

4.1 The Shimmer3 IMU Unit used in the wearable system. 34

4.2 The smart plate system showing the polymer plate mounted on top of the base station. The plate has three compartments for the serving of different food types. 35

4.3 Picture taken during a measurement of a meal in the IMU-REMY dataset. 39

4.4 Pictures taken during a measurement in the PLATE-REMY and PLATE-CHINA datasets. 40

5.1 The mastication muscles that move the lower jaw 47

5.2 The wearable capture setup: a Shimmer3 IMU mounted to a pair of glasses. Also shown is the local coordinate system of the sensor. 48

5.3 Internal structure of a differential capacitive MEMS accelerometer. 49

5.4 Example of a raw accelerometer signal recorded with the capture setup (top) and the same signal filtered with a high-pass filter (bottom). 51

5.5 Example accelerometer data showing the three dimensional accelerometer signal together with the binary annotation signal. 52

5.6 Average frequency spectrum of chewing versus not-chewing in the IMU-REMY dataset. This was obtained by averaging the spectra of all epochs of chewing and not-chewing respectively. 53

5.7 The operating principle of the segmentation process, illustrated with an example acceleration signal. 55

5.8 Effect of the window length on performance for the SVM classifier with quadratic kernel on the IMU-REMY dataset. 59

5.9 PPV-sensitivity curves for each of the evaluated classifiers validated using the IMU-REMY dataset. 59

5.10 PPV and sensitivity of the quadratic SVM for each participant in the IMU-REMY dataset. The dashed and dotted horizontal lines denote mean PPV and sensitivity over all participants, respectively. 60

5.11 Illustration of an accelerometer signal captured while chewing in a controlled environment. 64

6.1 Static force system with 1 mass. 70

6.2 Static force system with 2 masses. 71

6.3 Drawing of a strain gauge load cell. 73

6.4 The smart plate system with polymer plate and base station. . 75

6.5 Position of the load cells on the base platform. 75

6.6 Block diagram of the embedded capture system. 76
6.7 Differential amplifier and low pass filter for the load cell. 78
6.8 The custom PCB design for the smart plate. 78
6.9 Linearity test. 81
6.10 Drift of the system during 8 hours of constant load. 81
6.11 Repeatability of the system with a weight of 200g. Total deflection is 0.23 g. 82
6.12 Position test with a weight of 5 g. The locations where grid lines intersect indicate the true locations of the weight. 82

7.1 The test setup of the preliminary bite detection study. Served food in the lab (left) and home (right) environment is shown. The smart plate compartments are labelled 1 to 3. 85
7.2 Restricted eating versus unrestricted eating, captured with the smart plate system. 88
7.3 Hand drawn example of a single bite. 89
7.4 Hand drawn example of a partial bite. 90
7.5 Hand drawn examples of utensils resting and zero bite. 91
7.6 Plot showing the stable regions extracted from the raw total weight signal, denoted by black shaded areas. 93
7.7 The overall detection pipeline. 95
7.8 Hand drawn example of the classification results. Thick horizontal bars indicate a detection, thin vertical lines are the true labels. 96
7.9 Box plot showing the distribution of the weight error per compartment, relative to the total weight of the meal. Shown for both the PLATE-REMY and PLATE-CHINA datasets. 100

8.1 Bite parameters illustrated on a hand drawn weight signal. . . 107
8.2 Probability density function of the bite length for two subjects in PLATE-CHINA. It shows the distribution for a single meal per person. 107
8.3 Average bite parameters sorted by age in PLATE-CHINA. . . . 109

8.4 Total bite count per meal sorted by age in PLATE-CHINA. . . 109

8.5 Box plot showing the distribution of bite size for each person in the PLATE-REHAB dataset. 112

8.6 Box plot showing the distribution of bite length for each person in the PLATE-REHAB dataset. 113

8.7 Box plot showing the distribution of the time required to pick up a bite in the PLATE-REHAB dataset. 114

List of Tables

2.1 Example of a confusion matrix. 21

4.1 List of labelled activities in the IMU data. 36
4.2 Information about the IMU-REMY dataset. 39
4.3 Overview of consumed foods in the PLATE-REMY dataset. . . 42
4.4 Overview of consumed foods in the PLATE-CHINA dataset. . . 43

5.1 Overview of features f_j extracted from the acceleration signal. 56
5.2 The selection of features after applying the mRMR feature selection algorithm. Shown are the indices j that correspond to features in Table 5.1. 56
5.3 Validation results on the IMU-LAB and IMU-REMY datasets with 15 s windows. [mean±stdev.] 58

6.1 Mechanical and electrical characteristics of the CZL616C load cell. 74

7.1 The results of the preliminary bite detection tests. 87
7.2 Confusion Matrices: bite detection. 97
7.3 Confusion Matrix: bite localisation for the PLATE-CHINA dataset. 99

8.1 Participant characteristics of the PLATE-REHAB dataset. . . . 111

Chapter 1

Introduction

1.1 Malnutrition in the Ageing Population

Malnutrition is a frequent condition in the frailest of people in the ageing population [1, 2]. It is estimated that up to 15 % of community dwelling and housebound adults aged over 65 are malnourished, while 45 % are at risk [3, 4]. After institutionalisation, prevalence of malnutrition greatly increases. Up to 60 % of hospitalised older adults and up to 85 % of nursing home residents show signs of undernourishment [5]. There is a direct link between health status and malnutrition in older adults [6]. Malnutrition is associated with decreased muscle strength, poorly healing wounds, an increased hospital admission length and increased hospital mortality rate [7]. Furthermore, malnourished elderly are more prone to develop pressure ulcers and infections [8]. Preventing malnutrition by means of a targeted nutritional intervention could greatly improve the quality of life. Early recognition and treatment should therefore be included in the routine care of every older adult [3, 6, 9].

1.1.1 Self-reporting

Determining malnutrition can be done in a few ways. The first is by means of a self-report diary. These have been used to measure pain, sleep, illness or injury and health care use, as well as eating related issues such as binge eating, energy intake and expenditure in weight loss treatment [10]. In the case of malnutrition, the diary provides insight into two aspects of nutritional intake. The first is to monitor a person's eating behaviour and food consumption on a

daily basis in order to see if enough meals are consumed, and second, to record in detail all foods consumed for a nutrient analysis. The person is instructed to record all food intake, usually including location, time of day, quantity eaten, and nutrient values. A self-report diary is typically in paper-and-pencil format, but computerised solutions using a tablet-pc or terminal specifically catered to elderly people also exist [11]. It is clear, however, that a self-report diary has several limitations when used to self-monitor older adults. First and foremost, keeping track of food intake and the need to look up foods in a nutrient guide and record the amount of intake is a time consuming task. The self-monitoring protocol is seldom followed adequately, resulting in an incomplete diary [10]. Subjects may under or over report the food intake, either by accident or on purpose. Furthermore, limited literacy skills or bad handwriting also play an important role. Similar techniques such as 24-hour recalls, food records or food frequency questionnaires share the same limitations, especially in elderly care.

1.1.2 Nutritional Screening Tools

A different type and the most widespread tool for nutritional screening and assessment is the Mini Nutritional Assessment (MNA) [7]. The MNA contains 18 questions grouped into 4 parts: anthropometry, general status, dietary habits, and self-perceived health and nutrition states. Each question is graded and summed up to a total of 30 points. The result is defined by the following thresholds: a score below 17 indicates malnutrition; a score between 17 - 23.5 indicates a risk of malnutrition; scores above 24 indicates a good status.

Other tools such as the Geriatric Nutritional Risk Index (GNRI) [12] and Cumulative Illness Rating Scale (CIRS) [13] have also been used in combination with the MNA to provide further insight into the person's health status [3], [7].

While these questionnaires can be an effective tool in nutritional screening to prevent malnutrition, they typically have to be completed with the help of a care professional. Neither are they taken at routine intervals due to their time consuming nature. They are therefore not used as a preventative tool to detect malnutrition at an early stage. In case of housebound elderly receiving home care, tests such as the MNA are typically never administered unless ordered by a GP or only after admission to a hospital. The results of these tests are also not always on par with what caretakers observe on a day to day basis.

1.1.3 Automatic Food Intake Monitoring

A possible method to objectively measure food intake without relying on self-reporting tools is through the use of sensors that automatically capture eating activity. There are several methods to accomplish this. Wearable activity trackers can be used to detect movements of the body associated with eating activity. Chewing activity can be detected with sensors worn on the head or neck [14–18] while other body movements can be detected with wearables worn on the wrist or torso [19–21]. These methods have the advantage of being passive and offer high performance, at the cost of comfort. Depending on the type of sensor, this method may be suited for use in the elderly [16], but lack the ability to measure food quantity. Another approach is to use imaging techniques to detect and quantify types of food present on the plate [22–25]. The weight of the food can be estimated using 3D-reconstruction. The advantage of this approach is the low entry point if a smartphone is used to take pictures. The disadvantage is that manually taking pictures, often more than one per meal, is inconvenient. Furthermore, smartphone usage may prove difficult in the ageing population [26]. To accurately detect the amount of food consumed, weight sensors can be embedded in the kitchen or eating surface [27–30]. This, however, often requires an extensive adaptation to the existing infrastructure and measurements can only take place at the installation location.

1.2 Research Goals and Motivation

Traditional pen-and-paper based tools such as self-report diaries or questionnaires are the gold standard for intake monitoring, but are typically not suitable or simply not used when it comes to older adults. Sensor based automatic systems have been proposed in the literature as a way to ease the recording burden on the user, but existing research is targeted towards young to middle aged adults. Tools that are comfortable and convenient to use by older adults are virtually non-existent, while such tools could greatly lower the threshold for food intake monitoring and assist in the early recognition and prevention of malnutrition.

The goal of this thesis is to explore sensor based methods that are more suitable for use by older adults. Two modalities are presented: a wearable system and a weighing system in the form of a smart plate. The following research objectives are defined:

1. Investigate the needs and expectation of older adults, caregivers and other stakeholders regarding food intake monitoring through focus groups.

2. Explore the use of accelerometer based sensing for the detection of eating activity, in particular for the detection of chewing activity. Several sensors have been presented in the literature for the latter, but the accelerometer remains unused. The feasibility of the sensor is proven and tested on data recorded from older adults.

3. Develop a bite-detection and localisation algorithm for use with the smart plate during unrestricted eating. The algorithm is tested on data recorded from older adults. Furthermore, the use of bite-related parameters for the quantification of eating behaviour is explored.

1.3 Outline of the Thesis

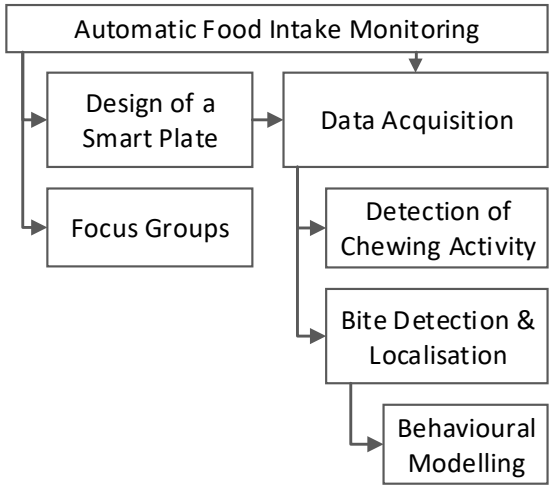


Figure 1.1: Schematic overview of the thesis structure.

Figure 1.1 shows a schematic overview of the thesis structure. The thesis presents two novel methods for the monitoring of food intake that are better suitable for use in the ageing population. To detect eating activity throughout the day, a wearable system with an accelerometer mounted on the person's glasses is used. This system detects the physical activity of chewing as an indicator of eating activity. The mechanical movement of the chewing activity results in vibrations which are transmitted via the skull to the frame of the glasses and finally to the accelerometer.

In order to also detect the amount of food consumed, a smart plate system is presented consisting of a base unit with pressure sensors and an off-the-shelf

polymer plate. Because all the hardware is contained in the base unit, the system is mobile and does not require any adaptation of existing eating surfaces. Furthermore, in addition to being able to measure the weight of individual bites, the location where each bite was taken from is estimated. In conjunction with the compartmentalised plate and with prior knowledge of which food was served in each compartment, this allows for an estimation of the total amount of ingested calories per meal.

Chapter 2 reviews existing works related to automatic food intake monitoring. The advantages and disadvantages are discussed to show that current methodologies are less suited for use by older adults. Furthermore, an introduction to machine learning techniques is given with an overview of the algorithms that are used in the thesis.

Chapter 3 gives an overview of challenges typically encountered when researching and developing technology for older adults. The *Engineers@Care Homes* project and insights gained from qualitative research (focus groups) during this project are presented.

Chapter 4 describes the data that is used in the thesis. Data was recorded in a lab, nursing home and hospital environment. This data is used for the training and evaluation of the machine learning models proposed in the thesis.

The wearable system with an accelerometer mounted on a pair of glasses is described in Chapter 5. The design of the smart plate system and bite localisation algorithm is outlined in Chapter 6, while the bite detection algorithm and complete detection pipeline, methodology and results are outlined in Chapter 7.

Chapter 8 introduces a preliminary study into using the data measured by the smart plate for behavioural modelling of patients. Parameters are extracted from the bites detected with the smart plate and used to show that there may be a link between these parameters and behavioural traits during eating.

Finally, a conclusion is formed in chapter 9. The presented work is summarised together with an overview of potential future work and its valorisation potential.

Chapter 2

Background

2.1 Introduction

Automatic food intake monitoring is the activity of measuring a person's food intake without requiring manual input from the user, typically using wearable or environmental sensors. There are two detection tasks associated with this problem. The first is the detection of eating activity by detecting body movements related to eating. For example, the physical movement of bringing food from a plate to the mouth or lifting up a glass to drink can be associated with eating activity. To establish a complete food diary, a measurement of the amount of food consumed is also required. This can be done using camera based systems or systems that measure the weight of the food consumed with pressure sensors. While several such methods for both tasks exist, they are often restricted to a controlled lab environment and share disadvantages that make them less suited for use by older persons. This chapter will give an overview of existing methods proposed in the literature. Furthermore, the detection of eating activity and food types is typically done with machine learning models. The relevant machine learning techniques that are used in the thesis are introduced.

2.2 Related Works

2.2.1 Wearable Systems

The majority of wearable systems focus on the Monitoring of Ingestive Behaviour (MIB), which is the measurement of periods of eating activity. MIB systems can help in better understanding eating behaviours and improve the accuracy of intake estimation. Automatically measuring periods of eating activity can aid in reducing the workload on the user, which in turn can reduce over or underestimation of the food intake [31]. To detect periods of eating activity, the two most popular methods are the measurement of chewing activity and detection of bodily movements related to eating activity.

A strain gauge can be taped to the lower jaw and used to measure characteristic movements of the jaw during chewing [18, 32–34]. This sensor is made from a piezoelectric film that converts the stretching of the skin into an electric signal. Sazonov and Fontana [18] used a Support Vector Machine (SVM) to classify epochs of chewing and non-chewing. Using a single strain gauge, they were able to obtain 81 % average detection accuracy. Measurements consisted of the meal and 20 minute pre and post sessions consisting of 10 minutes resting and 10 minutes of reading aloud. While performance is good, the downside of this approach is that it requires a wired sensor taped to the jaw, which makes it unsuitable for measurements over longer periods of time. Advances in miniature electronics should make it possible to produce a wireless variant of the sensor in the form of a patch, but this remains inconvenient as a method for regular and long term monitoring. In later work, Fontana *et al.* [19] improved on the single strain gauge by integrating it in a complete wearable system consisting of a hand-to-mouth proximity sensor worn on the wrist, an accelerometer module worn on a lanyard around the neck and the strain gauge fixed to the jaw, shown in Figure 2.1. They were able to achieve an average detection accuracy for eating activity of 90 % over the course of 24 hour measurement sessions. While this system offers a high performance, it requires multiple continuously worn devices in addition to the taped strain gauge, each requiring regular charging, heavily imposing on the comfort of the user.

A second class of MIB devices uses a wearable microphone in combination with audio processing techniques to detect characteristic sounds related to eating activity. Similar to the methodology of the strain gauge, a bone conducting microphone can be used to detect the sounds corresponding to chewing activity. The microphone can be built into a standalone device similar to a wireless ear bud [14, 20], or integrated in the existing hearing aid of the user [15, 35]. This method achieves similar performance as the previously discussed movement

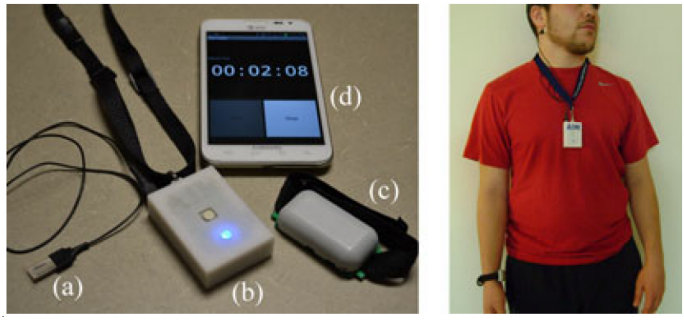


Figure 2.1: The Automatic Ingestion Monitor with (a) jaw strain gauge, (b) accelerometer module, (c) wrist mounted hand-to-mouth proximity sensor, (d) smartphone. Picture copyright: Fontana *et al.* [19].

based systems. When integrated into the existing hearing aid of the user, it offers the advantage of being easy to use with no additional discomfort. The disadvantage of using an audio-only system is that activities that do not produce sound during the meal cannot be included and ambient noise can impact performance. Performance can further be improved by combining both methodologies. Combining the jaw strain gauge with a microphone positioned over the throat to detect swallowing sounds allows for an increase in accuracy up to 94 % [33]. This, however, combines disadvantages of both systems. Furthermore, a throat microphone is very visible and can be stigmatising for the user.

2.2.2 Imaging Systems

A camera based system can be used to visually detect food intake from pictures or video recordings. Different food types can be automatically detected in pictures by comparing them to a pre-established database of food items [22, 36–39]. Yang *et al.* [39] present a classification scheme based on pairwise local feature mapping and detection with a SVM classifier. Pictures are taken from the side and require a white backdrop, making it unsuitable for use outside of a controlled environment. Wu and Yang [36] show a similar approach using video recordings of people eating fast-food. In more recent work, Convolutional Neural Networks (CNN) have been employed for this task. Ciocca *et al.* [37] demonstrated a CNN based feature extraction in combination with a SVM classifier. Using image segmentation techniques, different types of food can be detected in top-down pictures of food trays. Mezgec and Koroušić [38] show a

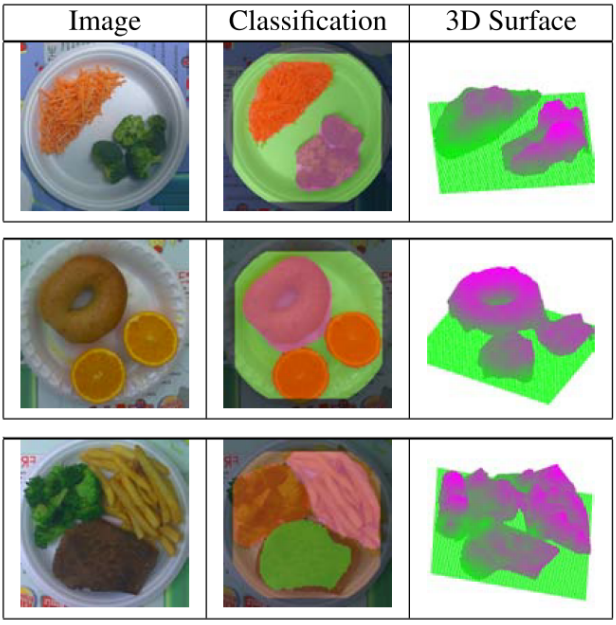


Figure 2.2: Image segmentation and volume estimation of different food types. The coloured areas in the Classification column indicate detected food types. Picture copyright: Puri *et al.* [25].

complete CNN pipeline for feature extraction and classification of pictures of food, but can only detect one type of food per picture.

Instead of manually taking pictures of the food or mounting a camera near the eating location, the camera can be worn on the head or body to mimic the natural field of view of the wearer. This is also referred to as egocentric video. Bolanos and Radeva [40] show that the type of food in front of the user (out of 101 food types) can be detected with an accuracy of 74 % using a body worn camera in free-living individuals. Jia *et al.* [41] show a similar approach with a button-style camera, but can only detect if food is present in the image and do not differentiate between food types. In recent work, Damen *et al.* [42] released the egocentric EPIC-KITCHENS dataset containing 55 hours of free-living activities related to cooking and eating food with extensive annotation of objects and food types in the image. Future work in this domain could lead to accurate detection of food types from egocentric camera images. The wearable camera requires little interaction once worn, but can be more complex in usage and requires more computational resources than other types of wearable device.

The detection of individual food types in a picture can be further augmented with 3D-reconstruction techniques to estimate food volumes [23–25, 43, 44]. By taking multiple pictures before and after the meal and subtracting the measured 3D volumes, the volume of consumed food can be estimated. While volume is not the same as weight and detecting calories from volume alone can be challenging, it can provide as estimation to the amount of food consumed. Puri *et al.* [25] show a complete classification and volume estimation pipeline. Using image segmentation and a SVM classifier, food types are first separated and classified according to food type. Three pictures taken from slightly different positions above the plate are then used to reconstruct a 3D volume of each food type. They were able to estimate the volume of the food on each plate with an average error of 5.8 %. The output of this process is illustrated in Figure 2.2. Pouladzadeh *et al.* [43] later followed a similar approach, but requiring only two pictures of the food: one from the top and one from the side. Using this technique, individual food types can be detected with an average accuracy of 92 % and an average volume estimation error of 6.4 %.

An advantage of camera based systems is that an already available smartphone camera can be used. The user still needs to manually perform the task of taking the picture, but this is already an improvement to manually entering the different food types for each meal. To estimate the amount of food consumed, however, the user needs to take multiple pictures from different angles before and after the meal, which can be inconvenient. Furthermore, the camera often needs to be calibrated before use. For older persons, the use of a smartphone can be problematic [26]. In this case a fixed camera set-up can be used, but this requires an extensive adaptation of the eating environment. Egocentric video can be a good alternative, but this method introduces new privacy concerns and can be stigmatising for the user.

2.2.3 Table Embedded Scale

Despite advances in wearable or imaging based ingestive behaviour monitors, weighing scales remain the gold standard for accurately measuring the amount of food for a given meal. As early as 1980, Kissileff *et al.* [45] proposed the integration of a scale into the existing eating surface to automatically measure the amount of food consumed during meals. The system, called the Universal Eating Monitoring (UEM), was comprised of a weighing surface with a single weight sensor connected to a computer. The user would place their meal on the weighing surface and the weight profile of the meal could be recorded as the meal was consumed. This allowed for the automatic measurement of the total weight of consumed food and calculation of cumulative intake curves that can provide researchers with detailed information on the effects of psychological, nutritional



Figure 2.3: Implementation of a table embedded scale. Food is served on a tray that is in turn placed on top of the weighing scale. The weight data from the scale can be used to detect individual bites during the meal. Picture copyright: Mattfeld *et al.* [29].

or pharmacological manipulations on changes in human appetite [46–48]. More recently, Mattfeld *et al.* [29] demonstrated a UEM integrated into cafeteria tables to measure individual bites and drinking events, illustrated in Figure 2.3. Detecting individual bites during the meal allows for a more in depth analysis of the microstructure of ingestive behaviour such as average bite size, time between bites and number of total bites [49, 50]. Chang *et al.* [28] show a dining table with multiple weighing surfaces to track multiple eaters simultaneously in a group dining setting. RFID tags were used to identify individual food containers on the table. Zhou *et al.* [30] propose a pressure textile matrix woven into a tablecloth in combination with a weighing surface. The pressure matrix allows for the tracking of different plates and activities on the plate, but their use of force sensitive resistors for the weighing element resulted in a poor measurement of meal weights. An alternative to embedding the scale in the eating surface, is using it in the kitchen where food is prepared. Chi *et al.* [27] presented a *smart kitchen* with weighing scales embedded in the counter top and under the stove to measure the amount of food being prepared for each meal.

While table embedded scales can offer a high weight accuracy, they require an extensive adaptation of the original eating surface. They are therefore often used only in a research environment. Furthermore, detecting individual bites during eating in an uncontrolled environment (e.g.: at home) is a difficult task with research in this domain remaining sparse.

2.3 Machine Learning Techniques

Several machine learning techniques are used in the thesis and are therefore introduced in this section. Machine learning is the science of teaching computers complex concepts and getting them to act without explicit programming instructions. The basic premise is to build an algorithm that can receive input data and use a mathematical model to predict or perform an action at the output.

The machine learning algorithms used in the thesis are of the *supervised learning* type. Supervised learners take example inputs and their known outputs and try to derive a generalised model from the data that maps inputs to outputs. Once the model is *trained*, an unknown input can be presented to the model to make a prediction about the input. These kind of models are popular in *classification* tasks, where the goal of the model is to classify unknown data into two or more classes. A popular example of a supervised classification algorithm is the spam filter, where emails (input) are classified as being spam or no spam (output). The model behind a spam filter has been trained with large amounts of known emails that have been labelled as spam or no spam. Once trained, it then automatically makes predictions about unknown emails without human interaction.

2.3.1 Data Representation

In a typical supervised learning task, data is represented as a table of *instances*. The raw data is seldom used directly for classification. Instead, each instance is described by a set of descriptors, or *features*, along with a label that contains its class. A feature is a measurable property of the instance that has been *extracted* from the data. Calculating features is often referred to as *feature extraction*. Choosing good features is crucial for the success of the algorithm. Features should be informative of the data, discriminating and independent. Features are typically stored in a feature matrix, with each row denoting an instance and each column a feature type. Class labels are stored in a column vector containing a class label for each row (instance) in the matrix.

2.3.2 Validation

A typical machine learning task requires at least two sets of instances: a training and test set. Both sets are derived from a known dataset with true labels assigned to each instance. The training set is used by the machine learning

model to learn the general concepts descriptive of the data. Once the model has been trained, it is used to make predictions on the test set. The true labels of the test set are not given to the algorithm during testing, only the instances are presented to the algorithm. The predicted outputs are then compared to the true labels of the test set to evaluate the performance of the model.

Typically, both the training and test set come from the same dataset that has been split. A large dataset may, for example, be split with a ratio of 70/30 for training and test set, respectively. For small datasets, where such a split would result in a small test set, cross-validation is often used. Cross-validation works by randomly partitioning the dataset into two sets and performing training and testing with each set. Multiple rounds are performed with different partitions and the classification results are averaged over the rounds. This allows the model to be tested with more data than would be possible with a traditional split, while still maintaining a large enough training set. It is a good way to estimate the generalisation of the model. A commonly used method is k-fold cross-validation, where the dataset is split into k equally sized subsets. One subset is retained for testing and the remaining $k-1$ sets are used for training. This is repeated k times, with each subset used for testing exactly once, and the results are averaged.

The method used in the thesis is leave-one-person-out cross-validation, which is a variant of k-fold cross-validation. Instead of randomly partitioning the dataset, at each round of cross-validation, data from 1 person is retained for testing and data from all other persons in the dataset is used for training. This is done for each person in the dataset and the results are averaged. This way, data from one person is never shared among the test and training set, which could bias the results.

2.3.3 Support Vector Machine

The Support Vector Machine (SVM) is used for binary classification between two classes. The SVM attempts to construct a hyperplane in multidimensional feature space that separates the two classes. When an unknown instance is presented to the SVM, a prediction is made based on which side of the hyperplane the instance is located [51–53].

Given a training set defined as $\{\vec{x}_i, y_i\}_{i=1}^{N}$, with $\vec{x}_i$ being the input feature vector and $y_i \in \{-1, +1\}$ the class labels, a linear hyperplane that separates both classes can be defined as:

$$\vec{w}.\varphi(\vec{x}) + b = 0 \tag{2.1}$$

with $\vec{w}$ an unknown weight factor, b a constant bias, and $\varphi(\vec{x})$ a mapping function that maps the d-dimensional input vector $\vec{x}$ into a d_h-dimensional space. For linear SVMs, the mapping function maps the input linearly to the output and the hyperplane retains the same dimensionality as the input vector. When separation by a non-linear hyperplane is required, instead of redefining the separation plane into a higher-dimensional space, the input vector is remapped with the mapping function instead. The hyperplane is kept linear, but the input vector is remapped to a higher dimension d_h so that linear separation can nevertheless be achieved. In practice, this is done with a kernel function.

Assuming that the hyperplane separates both classes perfectly, the hyperplane is further defined by:

$$\begin{aligned}(\vec{w}.\varphi(\vec{x}_i) + b \leq -1) \Rightarrow (y_i = -1), \forall i = 1, ..., N \\ (\vec{w}.\varphi(\vec{x}_i) + b \geq +1) \Rightarrow (y_i = +1), \forall i = 1, ..., N\end{aligned} \tag{2.2}$$

or equivalently:

$$y_i(\vec{w}.\varphi(\vec{x_i}) + b) \geq +1, \forall i = 1, ..., N \tag{2.3}$$

Given an unknown input vector $\vec{x}'$, the output prediction $\hat{y}$ of the SVM can thus be defined as:

$$\hat{y}(\vec{x}') = sign(\vec{w}.\varphi(\vec{x}') + b) \tag{2.4}$$

The distance D from a point (x_0, y_0) to a line $Ax + By + c = 0$ can be found by:

$$D = \frac{|Ax_0 + By_0 + c|}{\sqrt{A^2 + B^2}} \tag{2.5}$$

According to Formula 2.2, the instance $\vec{x}_i$ of class $y_i = +1$ that is closest to the hyperplane is defined by $\vec{w}.\varphi(\vec{x}_i) + b = +1$. Using Formula 2.5, the distance d_1 from this closest instance to the hyperplane can be calculated by:

$$\begin{aligned}d_1 &= \frac{|\vec{w}.\varphi(\vec{x}_i) + b|}{\sqrt{\vec{w}^2}} \\ &= \frac{1}{\|\vec{w}\|}\end{aligned} \tag{2.6}$$

Similarly, the distance d_2 of instance $\vec{x}_i$ of class $y_i = -1$ that is closest to the hyperplane is also given by $\frac{1}{\|\vec{w}\|}$. The total distance $d_1 + d_2$ is thus $\frac{2}{\|\vec{w}\|}$. During

the training phase, the SVM will attempt to maximise this distance between the hyperplane and the closest data points on both sides of the hyperplane. In order to maximise the distance $\frac{2}{\|\vec{w}\|}$, $\|\vec{w}\|$ needs to be minimised. This optimisation problem can be rewritten as:

$$\min_{\vec{w}} \frac{1}{2} \|\vec{w}\|^2 \quad \text{so that} \quad y_i(\vec{w}.\varphi(\vec{x_i}) + b) \geq +1, \forall i = 1, ..., N \tag{2.7}$$

This optimisation, however, assumes that both classes can be perfectly separated. It is also referred to as a *hard margin* SVM. In practice, there may be an overlap between instances that is no longer separable by the hyperplane. For this situation, a *soft margin* is introduced in the optimisation equation. This soft margin allows for a certain amount of instances to be placed on the 'wrong' side of the hyperplane during training. This is done by adding a *slack variable* ξ and cost parameter C, resulting in the following optimisation problem:

$$\begin{aligned} &\min_{\vec{w}} \frac{1}{2} \|\vec{w}\|^2 + C \sum_{i=1}^{N} \xi_i \\ &\text{so that} \rightarrow \begin{cases} y_i(\vec{w}.\varphi(\vec{x_i}) + b) \geq 1 - \xi_i, \forall i = 1, ..., N \\ \xi_i \geq 0, \forall i = 1, ..., N \end{cases} \end{aligned} \tag{2.8}$$

The slack variable ξ_i is assigned to each instance in the training set and can be thought of as the distance from the hyperplane if the instance is on the 'wrong' side, and thus result in a misclassification, and 0 otherwise. The term $\sum_{i=1}^{N} \xi_i$ is thus the sum of all errors. The cost parameter C is introduced to increase or decrease the influence of the errors on the optimisation problem. For very high values of C, the soft margin SVM is equivalent to a hard margin SVM and as little errors as possible are allowed in the model. For low values of C, the impact of errors on the optimisation is reduced and more misclassifications are allowed in the training set. The value of C should be carefully chosen for each classification problem. This is a so called *hyper-parameter*, which has to be defined prior to training and is not derived from the data itself.

2.3.4 k-Nearest Neighbours

The k-Nearest Neighbours (kNN) algorithm is a different kind of supervised learner that is popular in the literature. As opposed to the SVM, the kNN does not learn a model that is descriptive of the data. Instead, a prediction is made by evaluating a criterion directly on the feature data. Because of this,

the kNN is sometimes referred to as a *lazy learner*. Nevertheless, it can be a high performing classifier for certain classification problems [53, 54].

The training phase of a kNN consists only of storing the feature vectors and class labels of the training instances. To classify an unlabelled new instance, the majority label of k training samples closest to the new instance is assigned. The parameter k is a user defined constant that is typically tuned based on the distribution of features. To find the training samples that are closest to the new instance in feature space, a distance function is used. Euclidean distance is a commonly used distance function in kNN models, but different distance functions can be used depending on the properties of the dataset.

A kNN classifier is typically susceptible to noisy or irrelevant features. In noisy feature space, a high value of k will result in more classification errors, while a low value results in overfitting. Choosing a *good* value for k is thus important for the performance of the classifier. In supervised classification problems, k can be empirically determined by varying its value until the highest classification performance is achieved.

2.3.5 Decision Tree Learning

Classic Decision Tree

Tree models are popular in machine learning tasks because they are easy to interpret and can be represented graphically [53]. Decision trees can be used for a wide array of machine learning tasks, but we will only focus on the problem of supervised classification. A decision tree is essentially a rule model. The tree building algorithm attempts to learn a set of hierarchical rules that describe the relation between feature values and class labels. A decision tree is essentially a collection of 'if-then'-rules, but it is typically represented in the form of a linked tree as a set of nodes connected by branches. Each node in the tree represents a test for a feature. Depending on the result of the test, the algorithm may move down to the next connected node, or classify the feature as belonging to a class. Class labels are placed in the leaves of the trees, which are nodes that have no further connected nodes. To classify a new instance, the algorithm starts at the top of the tree and 'walks' down, evaluating the test at each node, until a leaf node is reached. The class label of the leaf node is then assigned to the instance.

Tree Bagging

Tree bagging, also called *bootstrap aggregation*, is an *ensemble learning* technique. The core concept of tree bagging is simple: instead of training a single decision tree, multiple (but different) decision trees are trained on the same training set. After training, a prediction is made by combining the outputs of all trees, typically using majority voting. Given a training set $\boldsymbol{X}$ and class labels $\boldsymbol{Y}$, the bagging algorithm selects with replacement a random sample $(\boldsymbol{X}_b, \boldsymbol{Y}_b)$ from $(\boldsymbol{X}, \boldsymbol{Y})$, for $b = 1, ..., B$, with B the desired number of trees. For each new set $(\boldsymbol{X}_b, \boldsymbol{Y}_b)$, a decision tree is trained and added to the bag of trees (also called an *ensemble*).

Tree bagging can provide better results on the same training set as compared to a single decision tree. Decision trees are sensitive to noise in the training data, but by using a bag of trees the noise can be averaged out, given that the trees are not correlated. By randomly sampling from the training set, the trees are not correlated because they are given different training sets.

Random Forest

The Random Forest (RF) classifier is a subset of tree bagging. The classic tree bagging approach samples on the training set and selects all features to be used in the training of a new tree. The RF classifier samples on the training set in the same way, but also employs *feature bagging*. At each split in the learning process, the RF trainer selects a random sample from the training set. From this random sample, a random set of features is also selected. For a training set with total number of features N, a random set of $\sqrt{N}$ features is commonly selected per split [55].

The goal of the RF classifier is to further decorrelate the decision trees in the ensemble. A single feature in the training set may be a very strong predictor for one class. In standard tree bagging, this feature will be selected in many of the trees in the ensemble. This causes the trees to be correlated, decreasing the performance of the ensemble, which can be prevented by using a RF classifier. RF classifiers can thus provide higher performance than traditional tree bagging.

2.3.6 Feature Selection

While machine learning models may be trained to give good separation on the training set, a model may not be generalised enough to classify other unknown instances. When a model fits the training data too closely without

generalisation for unknown instances, the model is said to be overfitted [56]. Overfitting can occur when the model is tuned too aggressively on noisy features. The model may generalise the noise or random fluctuations in the features. When an unknown instance without this noise is presented, the model will fail. Overfitting can also occur when there are too many features in the training set.

To prevent overfitting, a feature selection step is usually performed to prune noisy or bad features [57,58]. Other methods to prevent overfitting are regularisation, where a penalty is added to the features as model complexity increases [59], or by simply recording more data. There are several feature selection methods available in the literature. One of those methods is used in the thesis and introduced in this section.

Maximum Relevance Minimal Redundancy (mRMR)

Part of the feature selection step is choosing *good* features. There are several measures that constitute a *good* feature. Features should be descriptive of the data with regards to the class label, while features should also be complementary to each other. A feature selection that is commonly used to achieve this is the Maximum Relevance Minimal Redundancy (mRMR) method [60].

The mRMR algorithm ranks features according to mutual information with the class and maximises mutual information among themselves. Given two (possibly multi-dimensional) variables a and b, their mutual information I is calculated by:

$$I(a;b) = \sum_i \sum_j p(a_i, b_j) \log \frac{p(a_i, b_j)}{p(a_i)p(b_j)} \tag{2.9}$$

with p the probability density function.

Given a feature vector $\boldsymbol{X} = [\boldsymbol{X}_1, \boldsymbol{X}_2, ..., \boldsymbol{X}_N]$, class label vector $\boldsymbol{Y}$ and starting from an empty feature set $\boldsymbol{S}$, the first feature ($m = 1$) $\boldsymbol{S_1}$ is selected by maximising the mutual information with the class:

$$\boldsymbol{S}_1 = \{\ \boldsymbol{X}_i | \max\ I(\boldsymbol{X}_i; \boldsymbol{Y})\ \}, \forall i = 1, ..., N \tag{2.10}$$

Subsequent features $\boldsymbol{S}_m$ ($m > 1$) are then selected with the following formula:

$$\boldsymbol{S}_m = \left\{ \boldsymbol{X}_k \Big|_{\boldsymbol{X}_k \in \boldsymbol{X} - \boldsymbol{S}_{m-1}} \max \left(I(\boldsymbol{X}_k; \boldsymbol{Y}) - \frac{1}{m-1} \sum_{\boldsymbol{X}_l \in \boldsymbol{S}_{m-1}} I(\boldsymbol{X}_k; \boldsymbol{X}_l) \right) \right\} \quad (2.11)$$

If we run formula 2.11 for $m = 1, ..., N$, the selected feature set $\boldsymbol{S}$ will contain a ranked list of features ordered by maximum relevance and minimal redundancy. In practical applications, mRMR is not run for the entire feature set, but a stopping criterion is implemented to stop the selection at a certain number of features n . This stopping criterion depends on the machine learning problem and total amount of features. For small features sets, a simple method is to start with $n = 1$ and train and validate the classifier multiple times, each time adding $n = n + 1$ until performance starts to saturate or decline. This is also referred to as a *forward wrapper* feature selection. This method is utilised in the thesis.

2.3.7 Evaluation Metrics

Once a machine learning model has been trained and predictions have been made on the test set, the performance of the model can be evaluated. In a typical machine learning task, many different models may be trained with different parameters or features. An objective evaluation is then required to rank the models according to performance. This section gives an overview of the evaluation metrics utilised in the thesis.

Confusion Matrix

The performance of a classification model can be visualised in a *confusion matrix*. A confusion matrix is a table containing the number of true and false predictions for each class. Table 2.1 shows an example confusion matrix. Classified instances are placed in the matrix based on their predicted outcome relative to the real class value. For example, given a positive prediction $\mathbf{p}'$ for an instance with true class positive $\mathbf{p}$, the prediction is labelled as a True Positive (TP). The total number of TP is then shown in the confusion matrix. The remaining combinations of prediction versus real are False Negatives (FN), False Positives (FP) and True Negatives (TN).

		Predicted outcome	
		p′	**n′**
Real value	**p**	TP	FN
	n	FP	TN

Table 2.1: Example of a confusion matrix with TP = True Positive, TN = True Negative, FP = False Positive, FN = False Negative, p = positive class, n = negative class.

Performance Metrics

The true and false predictions in the confusion matrix can be used to calculate performance metrics. An overview of the performance metrics used in thesis are introduced here. These performance scores are typically denoted in percentages.

- Accuracy is the proportion of true classifications over all predictions in the dataset:
$$ACC = \frac{TP + TN}{TP + TN + FP + FN} \tag{2.12}$$
- Sensitivity or *recall* is the ratio of correctly detected positive samples to the total number of positive instances:
$$SENS = \frac{TP}{TP + FN} \tag{2.13}$$
- Specificity is the ratio of correctly detected negative samples to the total number of negative instances:
$$SPEC = \frac{TN}{TN + FP} \tag{2.14}$$
- Positive Predictive Value (PPV) or *precision* is the ratio of correctly detected positives to the to total number of instances detected as positive:
$$PPV = \frac{TP}{TP + FP} \tag{2.15}$$

- F_1 score is a measure that combines sensitivity and PPV. In real world machine learning tasks, a trade-off typically has to be made between both measures. The F_1 score can be used to maximise the combination of both measures:

$$F_1 = 2 \frac{SEN \ . \ PPV}{SEN + PPV} \tag{2.16}$$

2.4 Conclusion

This chapter has given an introduction to the machine learning algorithms used in the thesis and summarised the most relevant sensor based methods for automatic food intake monitoring. Existing methods in the literature are typically only suited for use in a lab environment due to their limited comfort when used for longer periods of time, or still depend on some form of manual data entry (e.g. manually taking pictures with a smartphone). This makes them less suitable for use by older adults over the age of 60. Relying on a manual data entry is one of the major factors contributing to mistakes in the monitoring protocol. Current wearable systems are often too complex to comfortably use on a daily basis. There is an acute need for automated monitoring technology that can measure food intake with limited to no user interaction.

Chapter 3

Food Intake Monitoring for Older Adults: Stakeholders, Needs & Requirements

3.1 Introduction

There are several challenges involved in the research and development of technology for older adults, which is the primary target group for the methods presented in this thesis. When designing such technology, it is important that the needs and requirements of all stakeholders are taken into account. This chapter gives an overview of those challenges and presents the results of three focus groups that were conducted to gain insights into the needs of doctors, older adults and caregivers.

3.2 The Engineers@Care Homes project

A number of research groups are already involved in technological research related to automatic food intake monitoring. Research focused on the use of this technology by older adults, however, is virtually non-existent. Furthermore, only a limited number of technological developments see a successful adaptation in a real-world scenario or translation into the healthcare market. A common cause for this limited translation is the technology-push approach that is typically

followed by academic researchers. The work is often conducted from the point of view of the researcher without keeping in mind the needs and desires of the relevant stakeholders [61, 62]. In the scope of this research, the stakeholders are the older persons, their family and health care professionals in charge of their care. Research into the needs of these stakeholders remains sparse, resulting in the researchers often using their own perception as a starting point for how the technological research should be conducted. By doing so, the implementation of the new technology may be ill-suited for use in their intended demographic [63, 64]. This often results in technology that improves on the state of the art and generates scientific publications for the research group, but has no chance of making it to market or being adopted for usage outside of a lab environment.

To prevent this from happening, researchers should follow a demand-driven, as opposed to a technology-push, approach. This requires a close collaboration between all stakeholders: healthcare providers, older adults and researchers. All stakeholders should be involved from the start of the research project. The advantage of this approach is twofold. Most importantly, input from each stakeholder can be used as a starting point for the research to make sure that there is a need for the technology, increasing the chance of a successful translation to market. Furthermore, a close collaboration where all stakeholders are already involved in the research process allows for easier and more extensive user testing. Developed technology should be tested on the intended target demographic, in this case older adults. Recruiting test subjects on an ad hoc basis without such collaboration can be a difficult process. A secondary advantage is that input from the stakeholders may introduce new insights previously unknown to the researchers, which can lead to new research tracks and collaboration opportunities.

To facilitate this kind of collaboration between the different stakeholders, the *Engineers@Care Homes* project was launched in January 2015 for a duration of two years. This was done by establishing an engineering lab in the OCMW (now Zorg Leuven) nursing home Edouard Remy in Leuven [65]. An engineer was present in the nursing home one day a week to collaborate with the nursing home residents and care providers during their research. To explore the knowledge transfer between the stakeholders and engineers, three technological test cases were selected to be completed in the nursing home. Two of these test cases touched on the problems of fall prevention and detection in the elderly. The third test case focused on the topic of this thesis: tackling malnutrition in the elderly through the use of automatic food intake monitoring.

Focus groups were organised with the relevant stakeholders to tailor the research questions and technology to their needs and adapt where necessary. Furthermore, the data used in this thesis was partly collected from nursing home residents.

By including the nursing staff early in the research process, the recruitment of test subjects was simplified. After going over the requirements, the staff typically already had a good idea about which persons would be suitable for testing, reducing the time needed to recruit test subjects.

A more detailed overview of the *Engineers@Care Homes* project and the fall detection/prevention related research can be found in the PhD thesis of my colleague, Dr. Greet Baldewijns, titled *Automatic Fall Risk Estimation and Improvement of Existing Fall Detection Techniques* (2016).

3.3 Stakeholder Study: Focus Groups

Gaining insight into the needs of stakeholders can be done through the use of focus groups. A focus group is a moderated session where relevant stakeholders are invited to engage in a discussion about topics that the researchers want to know more about. A focus group is typically conducted in the form of a Q&A session or interview, in which participants are asked about their perceptions, opinions, ideas or attitude towards a product or concept. As the name implies, sessions are held in a group setting to stimulate interaction between participants, which can lead to further insights into previously unknown problems.

3.3.1 Older Adults

The first major stakeholder is the older adult. A focus group was therefore organised with 8 older adults. During this session, we engaged in a discussion about existing wearable technology, activity trackers, smart home appliances and food intake monitoring tools. This gave us a better idea about how the older adults perceive this technology and to gauge their interests into different types of sensor systems. In the later stages of the research and technical implementation, care was taken to include this feedback to make sure that what we're developing would actually be used by older adults.

A selection of questions and topics that were discussed:

- Willingness to wear a wearable device (e.g. activity tracker, smart watch)
- What properties should a wearable device have?
- Are there any concerns or reservations about using a wearable device?
- Awareness of daily health and *quantified self*. In what way would you like to gain better insight into your health status?

- How involved are you in your nutritional health?
- Would you use a wearable device to track your food intake over time?

Older adults' perceptions of wearable technology and smart home applications are generally well studied [61, 62, 64, 66–69]. This focus group re-affirmed some of the already known issues, but also revealed some new insights. For wearable devices, the older adults were generally accepting towards using them, but were primarily worried about battery life. Naturally, this is a concern for users of all ages, but the rationale of the older adult differed from that of younger users. Given a smart watch, for example, several older adults remarked that they had the habit of sleeping with their watch on to have quick access to the time when they wake up during the night. Since charging of the device typically happens at night, this is a problem. The issue is therefore not so much with the charging of the device itself, but with the smart watch not fitting into the routine of keeping the watch on at night. When it comes to wrist mounted activity trackers, the older adults felt that the device should present more information directly accessible on the device, instead of having to use a smartphone to access the information. While some participants used tablet or desktop computers at home, none of them had a modern smartphone, which re-affirms that smartphone adaptation is low among elderly adults [70].

When it comes to malnutrition in the elderly, all participants were aware of this problem and understood the importance of early recognition. When asked if they would use a device that measures their food intake automatically, several participants remarked that they would not use this kind of technology of their own accord. Only if a doctor or family member explicitly told them that there was a problem and they had to undergo a monitoring period, would they consider using it. This, again, differs greatly from the perception of younger users, who are generally more interested in their health status and nutritional intake. The older adults were very interested in the concept of a table embedded scale or a smart plate, because these can be made to require very little input from the user. As long as food can be served in the same way the older adults are used to (routine is very important here, too), they wouldn't mind having one. One participant noted that they would be open to using some kind of smart plate, but a polymer plate would look out of place among their existing ceramic plates.

The older adult's family is another important stakeholder, but we did not hold a separate focus group with them. The older adults were questioned about involvement of family in the monitoring process. Several participants noted that they would not mind sharing food intake data with family members, but did not see a direct benefit of that. Overall, participants felt that this information would be best shared only with healthcare professionals.



Figure 3.1: Focus group with older adults about activity trackers and food intake monitoring technologies during the *Engineers@Care Homes* project.

3.3.2 Healthcare professionals

Besides the older adult, healthcare professionals such as doctors, nurses and caregivers are also stakeholders of an automatic food intake monitoring system. These stakeholders can be divided in two groups: medical and paramedical caregivers. The medical caregivers are doctors and nurses responsible for diagnosis and treatment of illness or disease, typically in a hospital environment. Paramedical caregivers, also called *allied healthcare professionals*, are those involved in the direct daily care of patients, distinct from nursing or medicine, typically through therapeutic or support services (e.g.: dieticians, occupational therapists, informal caregivers, etc.). The needs and requirements differ between both groups. Doctors and nurses are typically most interested in using the system as a tool to aid in the diagnostic process, dieticians can use it as a starting point for a better personalised diet, while informal caregivers could be interested in an early warning system to detect malnutrition. A focus group was therefore organised for each of the two groups of stakeholders.

Medical caregivers

A first focus group with medical care professionals was organised consisting of 2 general practitioners (GP), a head nurse (geriatrics), 4 general nurses and a nurse specialised in malnutrition. This time, participants were questioned about current practices regarding follow up of dietary intake in a hospital environment. Given a system that can automatically measure food intake in

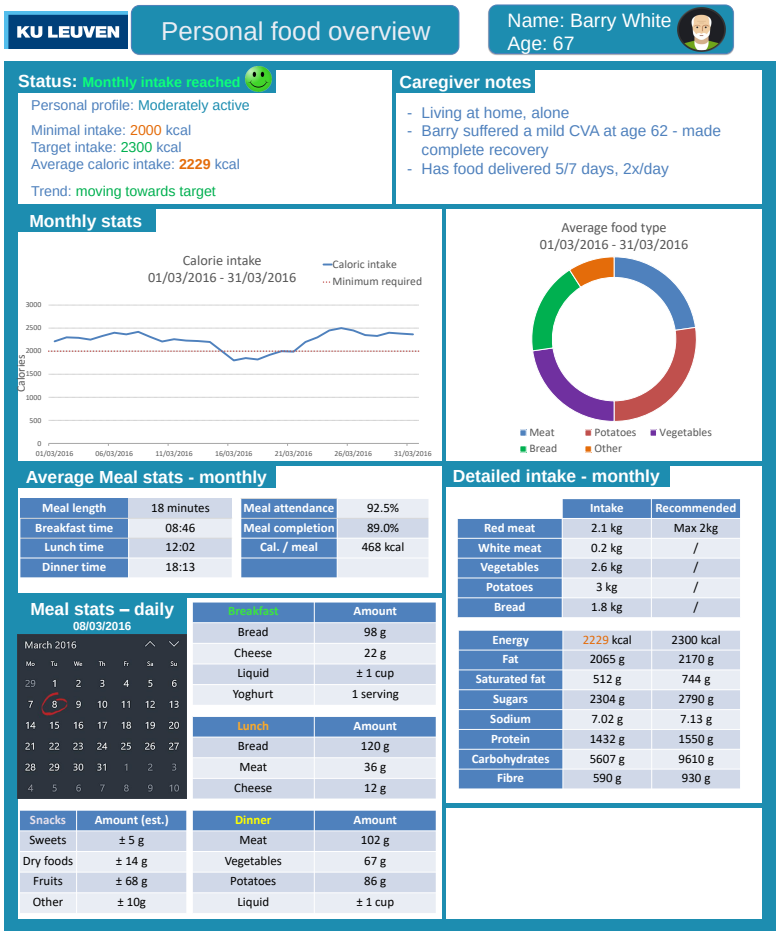


Figure 3.2: Mock-up of what a personalised food intake report could look like when presented to healthcare professionals.

patients, participants were also asked which parameters they would be the most interested in and how the output of such a system should be presented to them. They were presented with a mock-up of what the output of the system could look like in the form of a personalised patient report, shown in Figure 3.2. This report contains a detailed overview of the daily and monthly food intake that could be populated with data coming from the intake monitoring sensors. The information displayed in this report depends on the needs of the caregiver and situation in which it is used. Ultimately, it could be presented in an interactive format, where the healthcare professional can choose which data to display.

First and foremost, all participants confirmed the application potential for an automatic food intake monitoring system in a clinical environment. When intake monitoring is required for a patient, current practice consists of a visual inspection of the consumed food, followed by a manual data entry into the patient record. Automating this task through the use of sensors integrated in cutlery or the eating surface would reduce the nurse's workload. However, detailed information is often not required. A report such as the one we presented in the mock-up would be too detailed and an overview of ingested calories or weight alone would typically already be enough. This information is then used in the follow up of patients as they remain in the hospital, or as one of the parameters taken into consideration when forming a diagnosis or treatment plan. Furthermore, the GPs commented that they currently do not have access to longitudinal food intake data. Having historical food intake data, recorded over longer periods of time, readily available in the patient record could improve quality of care.

Paramedical caregivers

For the second focus group we invited paramedical caregivers involved in the care of older adults in the nursing home. This group consisted of a dietician, an occupational therapist and a speech therapist. The topics of discussion were similar to those used in the focus group with medical caregivers. While the context and reasoning for the use of a food intake monitoring system remains the same, the different needs of this group compared to the medical caregivers were confirmed. The dietician and speech therapist, for example, were interested in more detailed information than the doctors. A common problem with older adults is a decrease in chewing strength and problems with swallowing. Parameters such as chewing speed, bite size and time between bites were of interest to the speech therapist, who helps older adults with problems related to chewing or swallowing (dysphagia). Furthermore, the dietician was more interested in an overview of the different consumed food types and detailed contents of each meal over time, such as presented in the mock-up, whereas the doctors were generally only interested in more high level information such as overall calorie intake.

3.4 Ageing Population in China

While population ageing is a global phenomenon, China is the most populated country with the world's largest ageing population [71]. In 2013, over 202 million people in China were over the age of 60 and that number is expected

to nearly double to 402 million by 2040 [72, 73]. Moreover, life expectancy is projected to increase in the near future, further increasing the number of older adults with advanced age. China had 22.6 million people aged over 80 years old in 2013. This number is estimated to quadruple to 90.4 million by 2050, resulting in the world's largest population of this most advanced age group [74]. Similar to the Western world, malnutrition has been shown to be prevalent in hospitalised older adults in mainland China and Hong Kong [75, 76] and among community-living older adults [77].

An automatic food intake monitoring system could therefore also be useful in China. There are some cultural differences, however, that may require a different approach when designing and implementing such a system. Eating habits in China differ greatly from those in the West. Solid foods are consumed using chopsticks and dishes are typically shared between everyone present at the table. Food is thus not served in the form of a single meal per plate per person, as is customary in the West. Staple foods such as rice are typically served to each person individually, but main dishes are placed in the middle of the table and everyone is free to take bite-size portions. Even in individual meal settings, for example in a hospital, a meal consists of a selection of different dishes that are eaten over rice and are seldom mixed together.

To get a better idea of these cultural differences, part of the work presented in this thesis was done in collaboration with Fudan University in Shanghai, China. During a six-month research visit at the Center for Intelligent Medical Electronics (CIME) in 2017, the intake monitoring system presented in this thesis was evaluated on a group of Chinese older adults. Data was collected in the Department of Rehabilitation Medicine at Huashan hospital Jing'An Branch, affiliated to Fudan University.

3.5 Outline and Rationale of the Proposed Solution

Older adults are generally open to the use of wearable devices as long as there is little to no impact on comfort [68]. This was also affirmed by the older adults in our focus group. Furthermore, enthusiasm towards the device increases when it can be embedded in an article of clothing or in something they already wear, like a ring or hearing aid.

It is with this concept in mind that we propose the use of an accelerometer mounted on the glasses of the older adult to detect eating activity. The sensor can measure the vibrations corresponding to chewing activity and use this as a way to detect when the person is eating. While not all older adults wear glasses daily, it can be an alternative to existing methods and the hardware may be

embedded into other devices such as hearing aids. In this thesis, we present a prototype with the sensor attached to a regular glasses frame to show the feasibility of this approach. The hardware is simple enough that the sensor could later be embedded inside the frame. Glasses are typically taken off at night and placed in the same place (e.g.: on a night stand). A wireless charging pad could be used to charge the glasses at night. Processing can be done on the embedded hardware itself and the results can be sent to a server via Bluetooth or Wi-Fi at night. This thesis contributes to the modelling of the chewing detection algorithm. Embedding the sensor into the glasses and integrating it into a complete monitoring system can be the topic of future work. The hardware prototype and algorithm are discussed more elaborately in Chapter 5.

Table embedded scales have shown promising results for the measurement of food quantity, but require adaptation of the existing eating surface and can only be used at the installation site. Our focus groups confirmed the need for a system that can be used in a hospital environment where adaptation of eating surfaces is difficult. We present a standalone smart plate system as an alternative. A custom designed prototype is shown, consisting of an off-the-shelf polymer plate mounted on a base station. The base station contains the sensors and electronics. The system is completely standalone, battery powered and wireless. Hence it does not require any adaptation of the eating surface prior to usage and can easily be used in a hospital or home environment. Furthermore, it can be made to operate completely in the background, with the device waking up from sleep mode and recording the food intake when it is being used. This way, the older person does not need to interact with the system, which can further increase adherence. A novel method to detect individual bites during eating is presented. Furthermore, the system can estimate the location of each bite on the plate. This is calculated based on the movement of the centre of mass during eating and uses only the sensors located in the base station. Sensors do not need to be integrated into the plate itself. One base station can thus be used with multiple plates. The prototype hardware and bite detection algorithm of the smart plate system are outlined in Chapter 6 and 7, respectively.

Chapter 4

Data Acquisition

4.1 Introduction

This chapter gives an overview of the datasets collected in the thesis and the capture setup that was used. Data was collected from older adults using a customised acquisition system. This data is later used for the training and testing of the models used for chewing and bite detection in Chapters 5 and 7, respectively.

4.2 Capture Setup & Protocol

The two prototypes that are used for data collection are introduced in this section: a wearable system and a smart plate. A more elaborate description of each system's hardware and the captured sensor data is presented in their respective chapters.

4.2.1 Wearable System

The wearable system consists of an Inertial Measurement Unit (IMU) mounted to a pair of wearable glasses. The IMU is a self contained, battery powered, unit that has an accelerometer, gyroscope and magnetometer. Because IMU units are readily available on the market, we opted to use an off-the-shelf product rather than designing our own. The IMU used in this work is a Shimmer3 IMU,



Figure 4.1: The Shimmer3 IMU Unit used in the wearable system. Picture copyright: Shimmer (2018).

shown in Figure 4.1. An IMU is also worn on each wrist using a wrist strap. The data collected from the wrist-mounted accelerometers is not used in the thesis, but could be useful in future work. During measurements, sensor data is logged to the internal SD card and later downloaded onto a computer. The system is discussed more in depth in chapter 5.

4.2.2 Smart Plate

The smart plate consists of a custom embedded system built into a base station. The base station contains a Printed Circuit Board (PCB), battery and weight sensors. Food is served on a polymer plate with three compartments that is separate from the base station. The polymer plate is mounted onto the weight sensors in the base station prior to the meal. Once assembled, the smart plate closely resembles a regular plate, with the exception of a slightly larger base, as depicted in Figure 4.2. Similar to the wearable system, data from the weight sensors is logged to an SD card and downloaded to a computer after the meal. More details on the electrical and mechanical design of the smart plate can be found in Chapter 6.

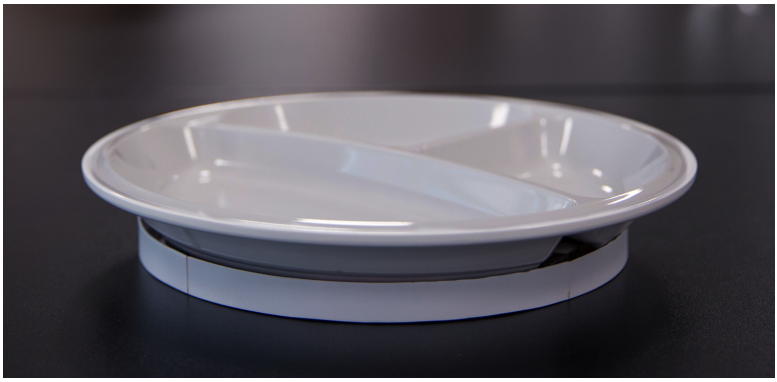


Figure 4.2: The smart plate system showing the polymer plate mounted on top of the base station. The plate has three compartments for the serving of different food types.

4.2.3 Measurement Protocol

While measurements were performed in a couple of different environments, the capture protocol remained the same for all measurements. Participants were asked to consume a meal while using the capture system. In the case of the wearable device, this meant wearing the glasses and wrist straps with the IMU units. Similarly, food was served on the smart plate and participants ate food from the plate like they were used to. Food was typically served from a central kitchen and came already served on a plate or tray. Foods were therefore transferred onto the smart plate prior to each meal. A researcher was always present during the measurement to assist with the capture setup. Measurements were manually started and stopped before and after the meal respectively. Data from the wearable system's and smart plate's SD cards was then transferred to a computer.

A different type of food was typically served in each compartment of the smart plate. Prior to the food being served on the plate, the food in each compartment was weighed with a kitchen scale. The remaining food from each compartment was again weighed after the meal. This serves as a reference measurement for the consumed food that is later used to compare the intake detection algorithm to.

Table 4.1: List of labelled activities in the IMU data.

Activity (start & stop time)
Chewing
Drinking
Cutting
Preparing food
Eating soup

4.2.4 Labelling

Participants were also video recorded for annotation purposes. Both the IMU unit and smart plate produce a one-dimensional sensor signal that does not contain any contextual information about what is happening at each point in time. An annotation process is thus done afterwards to label periods of time with the activities that took place during that time. After each measurement, an observer watched the video and labelled activities related to the food intake. Using a video image for labelling requires a synchronisation between the video capture system, the wearable device and smart plate. Synchronisation was achieved by referencing each system to the same clock. Video capture was done on a computer and at the start of each measurement, the date and time of the computer is loaded into each system at the same time. Both the IMU units and the smart plate contain a Real Time Clock (RTC) module that accurately keeps time. The computer time is shown in the video image so that each video frame can be referenced to an absolute point in time. Similarly, each sample recorded by the IMU and smart plate has a time stamp with the time from the RTC. Each frame in the video image can then be referenced to the corresponding data point from each sensor at the same point in time.

Table 4.1 gives an overview of the activities that were labelled in the IMU data. Each label contains a start and stop time. These are the moments when each activity is observed to begin and end, respectively. For the chewing activity, the labelling is started the moment a piece of food is brought to the mouth and the participant begins chewing. When the subject visibly stops chewing, the annotation is stopped. This typically happens shortly after the subject swallows the food, but some participants had the habit of swallowing parts of their food and continue the chewing motion, in which case the annotation is also continued.

The annotation of the smart plate data contains labels that correspond to individual bites taken from the plate. The moment when a participant takes a bite from the plate is labelled as well as the time when that bite is put in the mouth. The location of the bite on the plate is also recorded. Recall that the polymer plate has three different compartments. Each bite is thus also labelled with a compartment number that corresponds to one of the three physical compartments. Labelling of the smart plate data was done with the help of Dries Claes during his master thesis research.

4.2.5 Ethics

The sensor data used for the training and testing of the models presented in the thesis is anonymous. The raw signals captured by the IMU units and smart plate cannot be linked to an individual person. Each participant was assigned an anonymous identification number that is used to reference the data. The only personal information directly linked to this id number is the age of the participant. For administrative purposes, an excel sheet is kept that links person names to their corresponding id number, but this file is kept only by the researcher in an encrypted container and is never shared with third parties.

Furthermore, the video recordings of participants eating their meals are not used for any purpose other than labelling and visual inspection. The video recordings are not used in the models presented in the thesis. This was also communicated to participants and all participants gave their consent to use the video recording for this purpose.

Participants were given an info brochure containing all information relevant to the study as well as their rights. Prior to each measurement, the researcher explained the measurement protocol to each participant. All participants were informed of their right to stop the measurement and withdraw from the study at all times, which would not have any impact on their daily care. All participants signed an informed consent giving permission for both the capture of sensor data and the video recording for annotation purposes.

Ethical approval for the recording of the data was given by the relevant governing bodies. For the data collection in Belgium, ethical approval was provided by the Medical Ethics Committee of the Leuven University Hospital (ref. number S57883). Data collection in China was approved by the Huashan Hospital Institutional Review Board (ref. number KY2013-163).

4.3 Datasets

4.3.1 Belgium

Thee datasets were collected in Belgium. One was recorded in a lab environment (IMU-LAB) while the remaining two were captured in the nursing home (IMU-REMY and PLATE-REMY).

IMU-LAB

A small exploratory dataset was recorded with 5 young to middle aged adults wearing only the glasses IMU sensor. Participants were recruited among colleagues at KU Leuven Campus Geel. Each participant consumed a single meal and was asked to limit any activity other than consuming and chewing the food. Brief periods of talking with an observer are also included to simulate real world usage. This data was used to test the feasibility of the approach and if the glasses mounted IMU was able to measure chewing activity.

IMU-REMY

The results obtained with the IMU-LAB dataset confirmed that the IMU glasses were able to detect chewing activity. To further explore the usage of this system, a dataset was collected with older adults in an uncontrolled environment. Data was collected at nursing home Edouard Remy in Leuven. Participants were recruited among residents of the nursing home. The most healthy older adults were selected with good physical and cognitive abilities. Measurements of one meal per person were recorded from 13 older adults over the age of 65. This time, participants wore the full wearable system consisting of the glasses mounted IMU and wrist mounted IMU units on both wrists, as shown in Figure 4.3.

To get an idea of how the data would look like in a realistic environment, no restrictions were placed on participants when they consumed a meal. In a nursing home environment, residents typically have a fixed seat and meals are consumed in a social setting. For most measurements, participants were recorded in pairs of two to maintain this group setting. Participants were allowed to talk to each other, or stand up from their seats to get coffee, and do any activities that they would usually do while eating. Table 4.2 gives an overview of information related to the dataset.

Table 4.2: Information about the IMU-REMY dataset.

Number of participants	13
Participant age [mean ± stdev.]	78 ± 8
Total measurement time	409.3 min.
Total chewing time	93.8 min.
Avg. recording time per person [mean ± stdev.]	31.5 ± 6.5 min.
Number of labelled chewing epochs	153
Avg. duration per chewing epoch [mean ± stdev.]	37 ± 57 sec.*

Note * : Several participants had the tendency to chew almost continuously for the duration of the meal with little pauses between bites. This resulted in a high average time and standard deviation per chewing epoch.



Figure 4.3: Picture taken during a measurement in the IMU-REMY dataset.

PLATE-REMY

A second dataset was recorded at nursing home Edouard Remy in Leuven. Participants only used the smart plate system to consume a meal, as shown in Figure 4.4a. Data was recorded from 14 older adults recruited at the nursing home. One meal was recorded per participant. The average participant age and standard deviation is 82 ± 10. Similar to the IMU-REMY dataset, participants were asked to eat their meal like they were used to and no restrictions were placed on their actions. This was done to get a picture of what the sensor data looks like in realistic conditions. A total of 317 bite events were labelled. An overview of the different food types served to each participant is shown in Table 4.3. Also shown are the weights of consumed food for each compartment.

4.3.2 China

The final dataset used in the thesis was captured in China. Data was collected in the Department of Rehabilitation Medicine at Huashan hospital Jing'An Branch, affiliated to Fudan University. While participants were recruited from patients in the rehabilitation ward, the most healthy subjects with no upper body impairment and who had no issues with eating were selected. Several older family members were also recruited. A total of 18 older adults were recruited. Each participant consumed a single meal using both the wearable system and smart plate, as shown in Figure 4.4b. Average participant age and standard deviation is 63 ± 10. A total of 836 bite events were labelled. An overview of the different food types served to each participant is shown in Table 4.4. Also shown are the weights of consumed food for each of the three compartments. In the thesis, only the smart plate data is used and referenced as PLATE-CHINA.

Furthermore, an exploratory dataset was recorded from 6 patients who recently suffered a stroke and were undergoing rehabilitation. Patients in this group did have upper body impairment and had some kind of difficulty with eating. From this patient group, repeat measurements were taken. Over a period of 15 days, subjects consumed a meal with the smart plate system every 3 days, for a total of 5 measurements per person. This dataset is not used for the training of the bite detection model and as such, it is unlabelled. Instead, it is used to extract parameters in an attempt to quantify eating behaviour. The repeat measurements may allow for personal characteristics to be measured, and the upper limb impairment can be used as a behavioural characteristic that may be visible in the data. This dataset is referenced as PLATE-REHAB later in the thesis. This is further described in Chapter 8.



(a) PLATE-REMY (b) PLATE-CHINA

Figure 4.4: Pictures taken during a measurement in the PLATE-REMY (a) and PLATE-CHINA (b) datasets. Participants can be seen eating from the smart plate. A portable webcam and computer are used to record the video.

4.4 Conclusion

The datasets that were collected during the PhD research have been introduced in this chapter. All data was purposely collected for this research by the author. While there are public datasets available with accelerometer data, none contain a head-mounted accelerometer as proposed in this thesis. Data collected during eating is sparse, while data from older adults is non-existent. The data collected with the smart plate is similar to data collected by other researchers with table embedded scales, but this data is not publicly available. Furthermore, the smart plate is unique because it captures data from more than one weight sensor. Because of these reasons, our own data collection was required.

Table 4.3: Overview of consumed foods in the PLATE-REMY dataset.

id	Food type C1	C2	C3	Start weight [g] C1	C2	C3	End weight [g] C1	C2	C3	Total consumed [g]
1	fries	salad	goulash	109	70	197	66	24	74	212
2	mashed potatoes	mashed potatoes	meat balls	80	104	168	1	90	94	167
3	fried potato	fried potato	chicken fillet	80	101	131	16	35	80	181
4	salsify	mashed potatoes	pork	107	107	119	1	0	58	274
5	fries	chicory	meat stew	54	84	129	30	34	75	128
6	apple sauce	croquettes	turkey	142	113	93	0	78	5	265
7	cabbage	mashed potatoes	chicken fillet	71	83	94	71	83	20	74
8	mashed potatoes	mashed potatoes	cod	83	94	96	0	0	123	150
9	potatoes	peas	chicken fillet	88	66	144	40	50	20	188
10	potatoes	broccoli	salmon	62	73	80	27	21	27	140
11	croquettes	croquettes	meat stew	73	45	153	0	0	20	250
12	mashed potatoes	mashed potatoes	meat balls	69	89	70	3	85	12	128
13	-	-	spaghetti	-	-	283	-	-	8	275
14	mashed potatoes	mashed potatoes	fish stew	92	107	133	0	0	2	330

Note: C1, C2 and C3 stand for compartment 1, 2 and 3 respectively, which are the different compartments of the polymer plate. The start and end weights per compartment were weighed with a kitchen scale before and after the meal respectively.

Table 4.4: Overview of consumed foods in the PLATE-CHINA dataset.

id	Food type			Start weight [g]			End weight [g]			Total
	C1	C2	C3	C1	C2	C3	C1	C2	C3	consumed [g]
1	potatoes	spinach	rice	149	125	264	68	48	90	332
2	spinach	white fish	rice	132	67	224	82	0	7	334
3	white radish and pork	baked beans & fried egg	rice	72	73	153	10	0	0	288
4	pork	spinach	rice	87	93	208	0	0	0	388
5	mix. vegetables	chicken & beans	rice	76	119	325	0	69	0	451
6	tofu	chicken	rice	130	139	296	82	20	27	436
7	tofu	vegetable noodles	tofu	71	121	105	2	32	94	169
8	tofu & soy sprouts	mushrooms	rice	81	90	204	8	0	0	367
9	tofu & soy sprouts	pork	rice	120	98	227	72	47	80	246
10	beef	cabbage	rice	108	127	213	50	74	113	211
11	shrimp	spinach	rice	710	53	126	13	8	0	868
12	tofu & spinach	pork	rice	157	114	322	79	0	0	514
13	-	-	dumplings	-	-	284	-	-	31	253
14	pork	mix. vegetables	rice	126	116	142	67	79	29	209
15	spinach & tofu	rice	white fish	116	164	105	12	18	66	289
16	mix. vegetables	chicken	rice	110	81	153	1	1	0	342
17	broccoli	fried chicken	rice	119	111	165	43	16	17	319
18	white fish	tofu & spinach	rice	101	149	214	2	37	2	423

Note: C1, C2 and C3 stand for compartment 1, 2 and 3 respectively, which are the different compartments of the polymer plate. The start and end weights per compartment were weighed with a kitchen scale before and after the meal respectively.

Chapter 5

Detection of Chewing Using a Glasses Mounted Accelerometer

This chapter presents an alternative to existing wearable systems for the detection of eating activity. The accelerometer signal from an IMU mounted on a pair of glasses is used to detect periods of chewing activity. Data from older adults is used to train and validate a machine learning model to classify between chewing and not-chewing. In this chapter, the signal preprocessing, feature extraction, the detection algorithm and the results are discussed.

The work presented in this chapter has been published in the following papers and abstracts:

- Gert Mertes, Hans Hallez, Bart Vanrumste, and Tom Croonenborghs. Detection of chewing motion in the elderly using a glasses mounted accelerometer in a real-life environment. In *Proc. of the 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 4521–4524. IEEE, 2017.
- Gert Mertes, Tom Croonenborghs, Bart Vanrumste, and Hans Hallez. Towards detection of chewing motion in the elderly using a glasses mounted accelerometer. In *Proc. of the International Conference on Biomedical and Health Informatics (BHI)*. IEEE, 2017.

- Gert Mertes, Hans Hallez, Tom Croonenborghs, and Bart Vanrumste. Detection of chewing motion using a glasses mounted accelerometer towards monitoring of food intake events in the elderly. In *Proc. of the International Conference on Biomedical and Health Informatics (ICBHI)*, 2015.

5.1 Introduction

Previous studies have shown that eating activity can be detected by measuring characteristic movements of the body during eating. Several researchers have proposed measuring chewing activity as an indicator of eating. This can be achieved using a microphone worn in close proximity to the head [14, 15, 20] or a strain gauge taped to the jaw [18, 32–34]. This latter approach offers the highest performance, but it requires the sensor to be taped to the lower jaw of the subject. This approach is feasible in a research environment, but for daily or regular use by older adults it is inadequate.

Older adults are generally more accepting towards wearable devices if they can be embedded into something they are already wearing, such as clothing or a ring. An extra microphone embedded into a person's hearing aid can be used to measure the sound generated by chewing [15]. This could be a less intrusive alternative to the strain gauge. This method, however, is sensitive to ambient noise and cannot measure any other activities related to food intake.

A sensor typically used for activity recognition is the accelerometer. These are frequent in the literature for the detection of various human activities [78–84]. Accelerometers have also been used for the detection of eating or cooking activity, but only by detecting overall bodily movement related to this activity, such as a sensor worn on the wrist to detect arm motion related to bringing food to the mouth. At the time when the research in this thesis was performed, no research group had proposed the use of an accelerometer for the detection of chewing activity. Accelerometers have the advantage that they can be used for a wide array of activity detection tasks in addition to chewing detection, such as fall detection or recognition of other activities [85, 86].

Since accelerometers detect changes in acceleration, they need to be physically attached to the moving body of interest. In order to use an accelerometer to detect chewing activity, one could be attached to the lower jaw of a subject to measure the up and down movement of the jaw. This method has little advantage over the strain gauge, which works in a similar manner and has already been shown to work with high performance. An alternative method is proposed in the following sections.

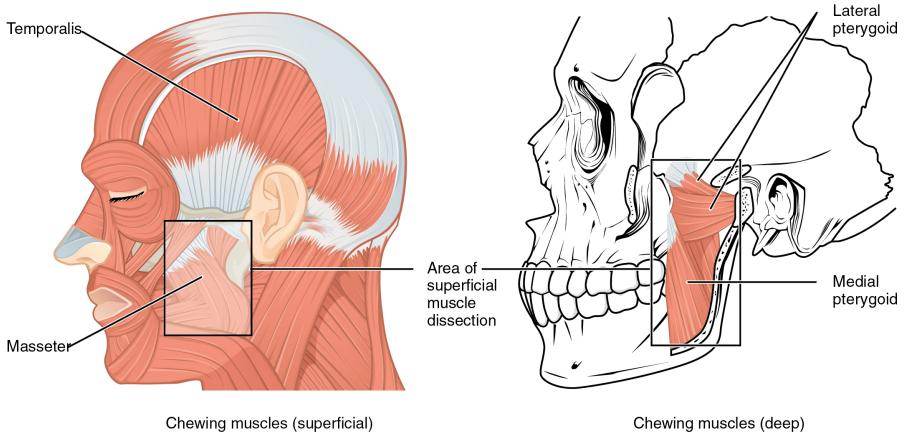


Figure 5.1: The mastication muscles that move the lower jaw [90].

5.1.1 Anatomy of Chewing

The chewing of solid food has been shown to generate vibrations that travel through the skull [87–89]. Furthermore, the mastication muscles around the ear, temple and lower jaw contract and move to some degree that is visible from the outside [90]. Opening and closing of the jaw is done primarily by the temporalis and masseter muscles, respectively. These are illustrated in Figure 5.1. These muscles lie superficially on the skull and their movement is visible on the outside. It is this observation that motivates us to capture this movement with an accelerometer. Directly attaching the sensor to the head has little advantage over existing techniques. We therefore propose the use of glasses as an intermediary platform to attach the sensor to. The ears of eyeglasses make direct contact with areas of skin that move as a result of mastication muscles contracting during chewing. Furthermore, the vibrations coming from the grinding of food can also be transmitted from the skull to the ears of the glasses, where in turn they can be picked up by the accelerometer.

5.2 Methods

Figure 5.2 shows the capture prototype with Shimmer3 IMU. A regular pair of glasses with removed lenses is used. To facilitate easy setup during the data collection and uniformity of the captured data, we opted to use the same pair of glasses for all test subjects. Because an off-the-shelf IMU unit is used,



Figure 5.2: The wearable capture setup: a Shimmer3 IMU mounted to a pair of glasses. Also shown is the local coordinate system of the sensor.

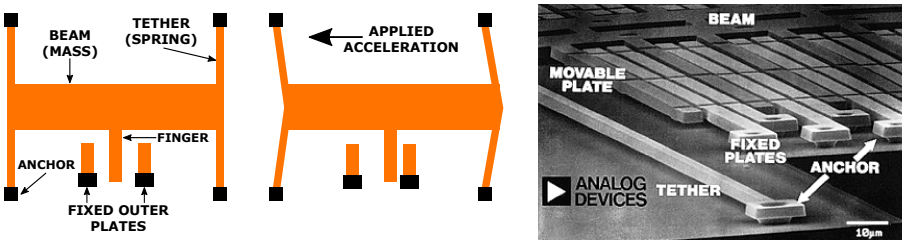
the capture setup is not very ergonomic. In future work, a custom integrated system could be made with a smaller form factor than the Shimmer3 IMU, or be fully integrated into the glasses. The purpose of this approach is to show the feasibility and test whether or not the sensor can capture chewing activity. While the capture setup looks bulky, it is comfortable to wear and the IMU is light enough to not be noticeable while wearing.

5.2.1 Accelerometer

The Shimmer3 IMU contains three sensors: an accelerometer, gyroscope and magnetometer. These measure acceleration, rotation and magnetic fields respectively. For the detection of chewing motion, only the accelerometer is used. This sensor can measure the vibrations from the chewing motion conducted through the ears of the glasses.

Accelerometer chips are micro-electro-mechanical systems (MEMS), typically made from a silicon substrate. An inertial mass is allowed to deflect or move under acceleration. A sensing element measures the displacement of the mass, which is proportionate to the acceleration. There are a number of different sensing methods, the sensor in the Shimmer3 is of the differential capacitive type. The inertial mass is surrounded by stationary electrodes by a narrow gap. When the mass moves, the capacitance between the mass and the electrodes changes. This change in capacitance is measured and converted into an acceleration signal by an integrated circuit [91].

Figure 5.3a shows a schematic of a finger type microstructure with a beam type mass. A finger is connected to the mass and positioned between the electrode



(a) (b)

Figure 5.3: Internal structure of a differential capacitive MEMS accelerometer.

plates. As the mass deflects under acceleration, the capacitance between the finger and the plates changes. Capacitance increases on one side and decreases on the other side between the finger and the plates. The differential signal is then measured and converted into an acceleration signal. The measured capacitance is proportionate to the overlapping area between the finger and the fixed outer plates. Because the fingers are typically in the order of around 100 µm, the capacitance of a single finger and plate pair is very small, resulting in a small resolution and high noise.

Increasing resolution and decreasing noise can be done by increasing the differential capacitance of the structure. This can be done by increasing the length of the fingers, but this can be difficult and expensive. Instead, a multitude of fingers and plates are positioned in a comb structure and electrically connected, increasing the differential capacitance [92]. Figure 5.3b shows a picture taken with an electron microscope of the comb structure. The fingers and plates can be seen suspended from the substrate by anchors. The flexible tether is connected to the beam, allowing it to deflect parallel to the base substrate. In a typical accelerometer, three of such structures are present to measure acceleration in three dimensions.

Accelerometers measure so called *proper acceleration*, which is the acceleration experienced by a body in its own rest frame. It can be described as the acceleration relative to free fall that is experienced by all people and objects. In practice, this means that an accelerometer at rest will measure an acceleration of $1g$ pointing *upwards* due to earth's gravity, with $g \approx 9.81 m/s^2$. In contrast, an accelerometer in free fall, falling towards the earth, will measure an acceleration of 0. In the frequency spectrum, earth's gravitational pull is a DC component.

The sensor used in the Shimmer3 is a tri-axial accelerometer that measures acceleration in a Cartesian coordinate system. The total acceleration measured by the sensor is thus split up in three components: X, Y and Z. The local

coordinate system of the Shimmer3 IMU is shown in Figure 5.2. When laid flat on a table with the Z axis pointing upwards, for example, the sensor will measure a constant $+1g$ in the Z axis and zero in the X and Y axes.

To obtain the acceleration due to movement relative to earth, the gravitational component should be subtracted from the signal. In practice, this is typically done by filtering the signal with a high-pass filter to discard the DC component. For some applications, the DC component can be of interest and the opposite can be done. The relative acceleration can be filtered out with a low-pass filter to retain only the DC component, which contains information about the sensor's orientation relative to earth.

An example signal, captured with the capture setup from Figure 5.2, is shown in Figure 5.4. Only the X-axis is shown. The DC offset visible in the raw unfiltered signal is due to earth's gravity. This is to be expected because when worn, the X-axis of the accelerometer is pointing down almost perpendicular to earth. Because the positive direction of the axis is pointed downwards, and the gravitational component is perceived as upwards by the sensor, the signal has an offset of approximately $-9.81 m/s^2$. When filtered with a high-pass filter, the DC component is removed and the signal is centred around zero.

5.2.2 Data

Data was collected in both a controlled and uncontrolled environment. The IMU-LAB dataset was recorded in a lab environment with 5 young to middle aged volunteers recruited on campus. This dataset is used for an exploratory study to determine the feasibility of the approach. The IMU-REMY dataset was recorded in a nursing home with 13 healthy older adults. This dataset is used to study the influence of real-world eating conditions on the captured data. From this dataset, two participants are omitted. One participant suffered from Parkinson's disease with upper body tremors. The acceleration signals were entirely predominated by the tremor and no chewing activity could be detected. Another participant was given a personalised meal consisting of soft foods, resulting in no chewing activity for this person. Both subjects were excluded, resulting in a dataset of 11 older adults. In both datasets, participants consumed a single meal while wearing the capture setup. A more detailed overview of both datasets is given in Chapter 4. For all measurements and datasets used in the thesis, the accelerometer signal is sampled at a sample rate of $f_s = 128\,\text{Hz}$.

Figure 5.5 shows an example segment of data from the IMU-LAB dataset. Shown are the three axes of the accelerometer together with the annotation signal for the chewing activity. The high level of the annotation signal indicates that chewing is taking place during that time. There is a visible difference

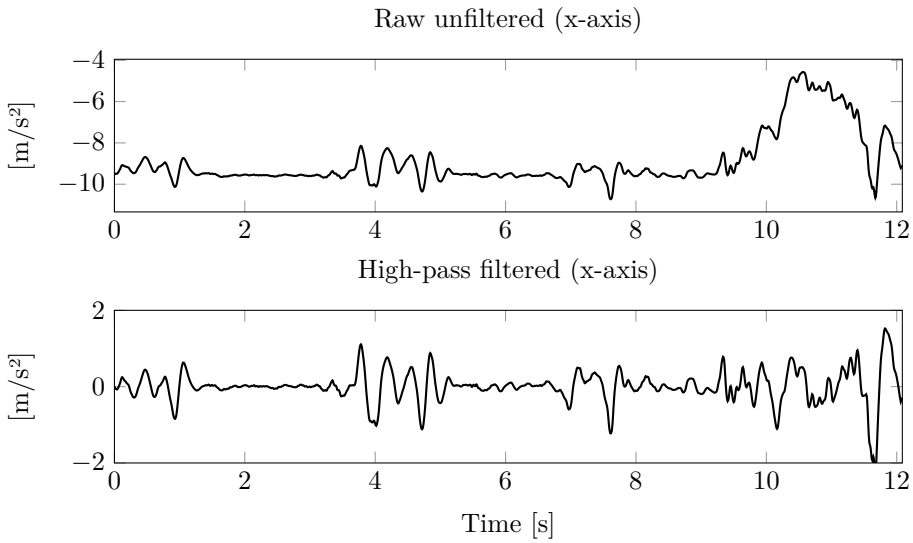


Figure 5.4: Example of a raw accelerometer signal recorded with the capture setup (top) and the same signal filtered with a high-pass filter (bottom). Shown is the x-axis only, which is pointed perpendicular down to earth when the glasses are worn. The gravitational component of $g \approx -9.81 m/s^2$ is visible in the unfiltered signal.

between periods of chewing and not-chewing in the accelerometer signal. When chewing occurs, the energy of the signal increases and high intensity peaks can be observed. Verification with the video data shows that these peaks correspond to the closing of the jaw when biting down on solid food.

Figure 5.6 shows the average frequency spectra of epochs of chewing and not-chewing from data in the IMU-REMY dataset. This was obtained by taking the average of the spectra during epochs of chewing and not-chewing respectively. The chewing activity contains a distinct peak in the 10-20 Hz range over the spectrum of not-chewing. Not-chewing also contains more low-frequency energy as opposed to chewing.

This difference in frequency spectra gives an indication that features extracted from the frequency domain may be good descriptors of both classes. Because the difference in chewing and not-chewing is also visible in the acceleration signal, features may also be extracted from the time domain. We therefore propose a mixed domain feature extraction. This is further described in the remaining sections.

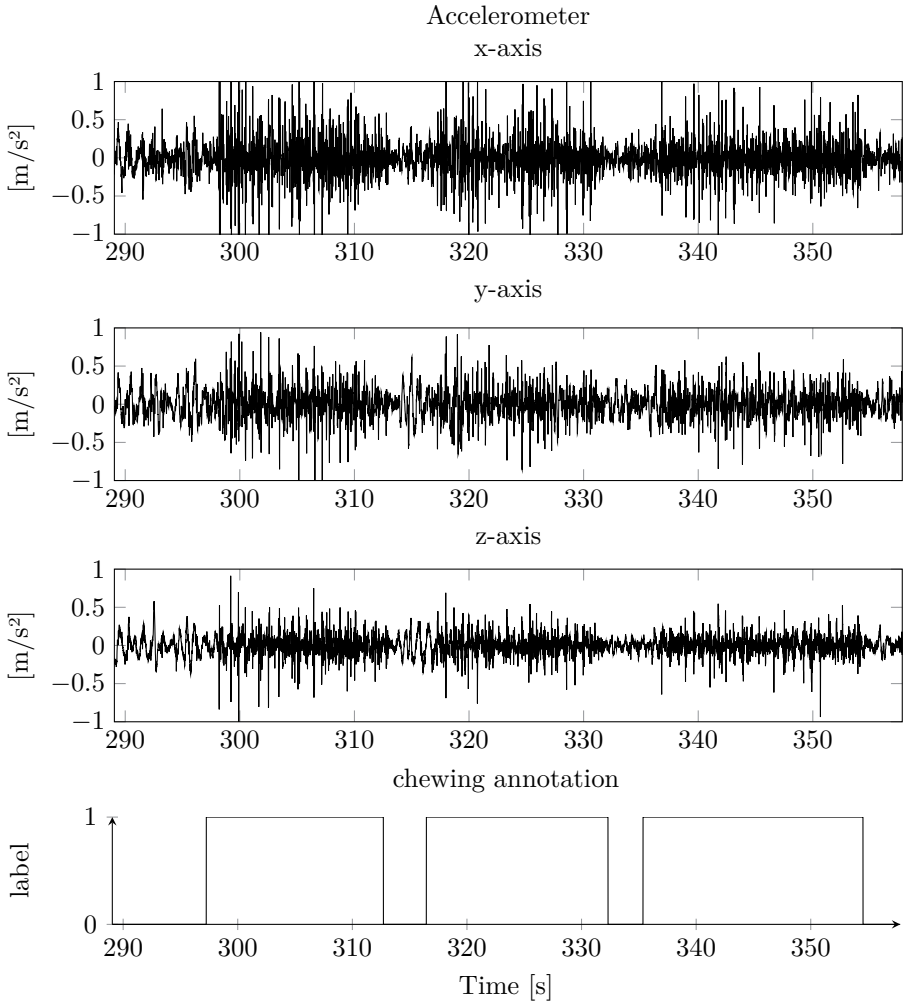


Figure 5.5: Example accelerometer data showing the three dimensional accelerometer signal together with the binary annotation signal. A 1 in the annotation signal indicates chewing activity. The difference between chewing and not chewing is visible in the raw accelerometer signal. Periods of chewing have higher peaks and more energy compared to not chewing.

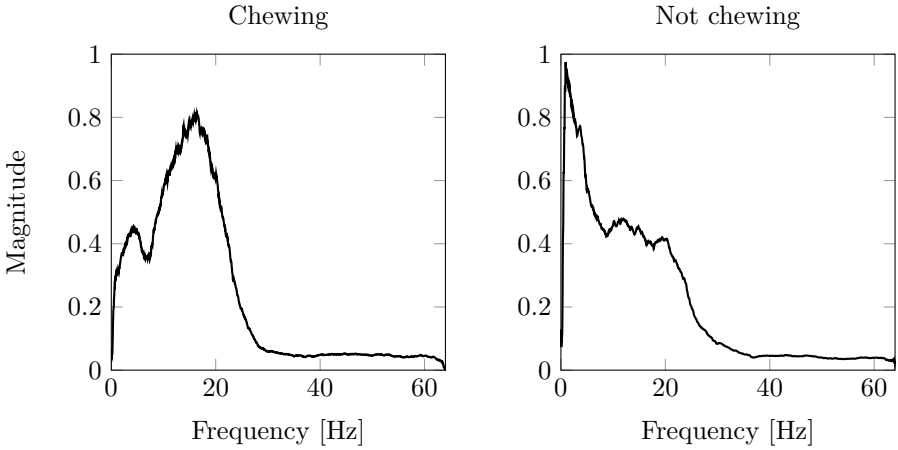


Figure 5.6: Average frequency spectrum of chewing versus not-chewing in the IMU-REMY dataset. This was obtained by averaging the spectra of all epochs of chewing and not-chewing respectively.

5.2.3 Pre-processing

The three axes of the accelerometer signal are hereby denoted by ACC_x, ACC_y and ACC_z. The raw accelerometer signals are each filtered with a high-pass filter to discard the DC-component. Recall that this component is due to earth's gravitational pull. While this component does contain information about the orientation of the sensor relative to earth, for the application of measuring vibrations due to chewing, this component can be discarded. This centres the signal around zero, which also simplifies the feature extraction later on. The filter used is a linear phase FIR filter with cut-off frequency $f_c = 0.5Hz$.

From the three (unfiltered) acceleration signals, the net acceleration r is calculated with following formula:

$$r = \sqrt{ACC_x^2 + ACC_y^2 + ACC_z^2} \tag{5.1}$$

This is the magnitude of the total acceleration experienced by the sensor. It is typically used in situations where the orientation of the sensor is unknown as it is uninfluenced by sensor rotation. For our application, the orientation of the sensor is known, but r can nevertheless be used in the feature extraction to produce features that may be of interest. The net acceleration r is subsequently also filtered with the same high-pass filter.

5.2.4 Feature Extraction

Machine learning algorithms work with *instances* that contain *features* extracted from the data. As described in Section 2.3.1, a feature is a measurable property that can be used as a descriptor of the instance. For example, given an acceleration signal of a certain length, a simple feature that can be calculated is the standard deviation. For noisy signals with high accelerations, the standard deviation will produce larger numerical values compared to calm and clean signals. Standard deviation could thus be used to differentiate between busy and calm acceleration signals. An observant reader will note, however, that this kind of feature only produces *one* numerical value for the *entire* acceleration signal. Indeed, using this type of feature we can only make a prediction about the signal in its entirety. This poses a problem if we want to detect activities *inside* the acceleration signal, which is not possible by calculating features from the entire accelerometer signal.

To solve this problem, the acceleration signal is segmented into smaller sized chunks, also called *windows*. Features are extracted from each individual window and a class label is assigned to it. Each window represents an instance to be classified by the machine learning algorithm. This methodology is similar to how humans would visually recognise patterns in an acceleration signal. We look at small portions of the signal and then decide whether or not that window contains any features we associate with the pattern that we are looking for.

This process is further illustrated in Figure 5.7 on an example acceleration signal. Features are extracted from each window and stored in a feature vector $\boldsymbol{X}_i$. Given a number of windows M, a single feature f_j and the number of different features N, a feature vector of size $1 \times N$ can be written as $\boldsymbol{X}_i = [f_1, f_2, ..., f_N]$ for $i = 1, ..., M$. The total feature matrix of size $M \times N$ can then be defined by $\boldsymbol{X} = [\boldsymbol{X}_1, \boldsymbol{X}_2, ..., \boldsymbol{X}_M]^{\mathrm{T}}$. The class labels corresponding to each window y_i are stored in a column vector $\boldsymbol{Y} = [y_1, y_2, ..., y_M]^{\mathrm{T}}$.

Window size is an important parameter that is typically chosen based on the signal type, the kind of activity to be detected and the type of features. A too large window may contain data from several different activities, while a too small window may not produce features that are descriptive enough to classify the activity. This trade-off is typically empirically studied by varying the window size and choosing the value that produces the best classification performance.

Table 5.1 gives on overview of the extracted features. This selection of features is based on those that have shown promising results when applied to the detection of motion related to eating activity and general human Activities of Daily Living (ADL) [18–20, 32–34, 78–83].

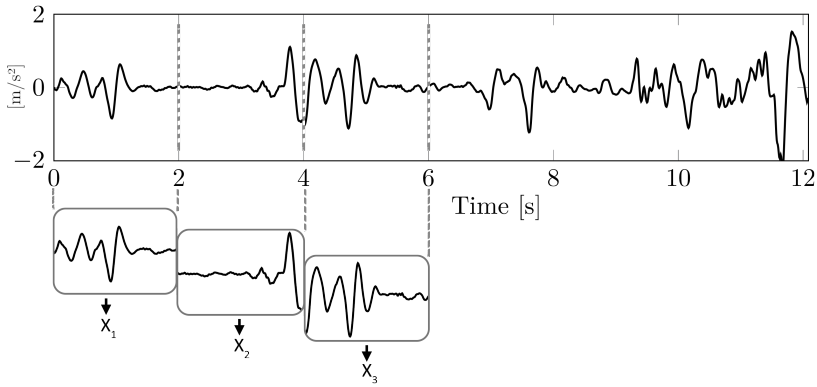


Figure 5.7: The operating principle of the segmentation process, illustrated with an example acceleration signal. Only the first three extracted windows are shown. In a real machine learning task, the entire signal is typically segmented.

After feature extraction, the final step before training is assigning a true class label y_i to each window. Class labels are assigned based on the labelling done by a human observer, using a video recording of the measurement. More details about this procedure are outlined in Section 4.2. Because the segmentation of the data is done arbitrarily with a fixed window length, there is a possibility that a single window contains data from multiple activities. When this is the case, a choice is made: the class label that is present for more than 50 % of the length of the window is assigned.

5.2.5 Feature Selection

Training the machine learning model with all extracted features will most likely result in overfitting [56]. An overfitted machine learning model will separate instances in the training set very well, but will fail to classify unknown instances. The extracted features may not all be descriptive of the data while some features may be redundant, which will have a negative effect on performance. A feature selection step is thus performed to remove redundant features.

A forward feature selection based on maximum relevance minimal redundancy (mRMR) is applied to the feature matrix, as outlined in Section 2.3.6. Two classes are taken into account: chewing and not-chewing. A *forward wrapper* feature selection is performed. The stopping criterion of the mRMR is set to $n = 1$ and the classifiers are repeatedly trained and validated while adding $n = n + 1$ until performance saturates or decreases. Table 5.2 shows the list of selected features. Shown are the indices j of features listed in Table 5.1.

Table 5.1: Overview of features f_j extracted from the acceleration signal.

j	Description (f_j)	j	Description (f_j)
1	SD of ACC_x (SD_x)	13	SD of ACC_y (SD_y)
2	SD of ACC_z (SD_z)	14	Power of ACC_x
3	Power of ACC_y	15	Power of ACC_z
4	# zero crossings of r	16	# zero crossings of ACC_x
5	# zero crossings of ACC_y	17	# zero crossings of ACC_z
6	# zero crossings of $\frac{d^2r}{dt^2}$	18	# zero crossings of $\frac{d^2ACC_x}{dt^2}$
7	# zero crossings of $\frac{d^2ACC_y}{dt^2}$	19	# zero crossings of $\frac{d^2ACC_z}{dt^2}$
8	{25, 50, 75} percentile of ACC_x	20	{25, 50, 75} percentile of ACC_y
9	{25, 50, 75} percentile of ACC_z	21	MAV of ACC_x (MAV_x)
10	MAV of ACC_y (MAV_y)	22	MAV of ACC_z (MAV_z)
11	Mean of {SD_x, SD_y, SD_z}	23	Mean of {MAV_x, MAV_y, MAV_z}
12	Dominant frequency of FFT(r)	24	{25, 50, 75} percentile of FFT(r)

Notes:
SD : Standard Deviation
FFT : Fast Fourier Transform
MAV : Mean Absolute Value, defined by:

$$MAV = \frac{1}{K}\sum_{i=1}^{K}|x_i| \quad \text{where } x_i \text{ is the } i\text{th sample in the window with length K}$$

Table 5.2: The selection of features after applying the mRMR feature selection algorithm. Shown are the indices j that correspond to features in Table 5.1.

Selected features	
4	17
6	22
9 [a]	24 [b]
12	

Notes:
[a] : only 25$^{\text{th}}$ and 75$^{\text{th}}$ percentile is selected
[b] : 50$^{\text{th}}$ percentile is selected

5.2.6 Training and Validation

Three different types of machine learning models are evaluated: a Support Vector Machine (SVM), k-Nearest Neighbours (kNN) and a Random Forest (RF). The SVM is evaluated with three different types of kernels: linear, Gaussian and 2^{nd} order polynomial (quadratic). For the kNN classifier, the number of k is evaluated between 1 and 10 and Euclidean distance is used as a distance metric. The RF is constructed with 50 bags of decision trees. The inner workings of these algorithms are outlined in Section 2.3.

All classifiers perform a binary classification between *chewing* and *not-chewing*. Chewing is considered the positive class, while not-chewing is the negative class. The not-chewing label contains all other activities during the meal such as sitting idle, preparing food, etc, but it also contains some eating-related activities such as drinking and eating soup. Since the goal of this study is to detect chewing activity alone, all other activities were pooled together in the general not-chewing class. Another reason why the other activities were not included as separate class labels is due to an unbalance in the amount of training samples of this class. Other eating related activities such as drinking are very sparse in the data and not enough examples are present in the data to make a generalised model.

Evaluation of the classifiers is done using a leave-one-person-out (LOPO) cross-validation scheme. Data of one person is used to evaluate a classifier trained with data from all other persons. This is repeated for each person in the dataset. The performance of the classifier is then calculated by averaging the performance scores over all subjects. The LOPO validation is done for each machine learning algorithm and dataset. It is done on each dataset individually, data is not shared between the IMU-LAB and IMU-REMY dataset.

In the IMU-REMY dataset, there are approximately 3 times as much segments of not-chewing than chewing. Due to this unbalance, PPV (precision) and sensitivity (recall) are used as performance metrics. In unbalanced datasets, accuracy is typically a bad indicator of performance because the unbalanced observations can bias the results [93, 94]. PPV and sensitivity are used because we are most interested in the positive class (chewing). Both measures say something about the positive predictions of the classifier. PPV is a good measure to evaluate how many positive detections were correct, while sensitivity gives an indication of how many chewing instances the classifier missed. Finally, F1 score is calculated as the weighed average of PPV and sensitivity. This is used to select the classifier with the highest combination of PPV and sensitivity.

Table 5.3: Validation results on the IMU-LAB and IMU-REMY datasets with 15 s windows. [mean±stdev.]

	IMU-LAB		
Classifier	PPV [%]	Sensitivity [%]	F1-score
RF	86.6 ± 12.0	85.3 ± 13.0	0.86
kNN	84.6 ± 13.7	86.2 ± 18.8	0.85
SVM - linear	87.8 ± 15.8	85.7 ± 19.3	0.87
SVM - gaussian	85.7 ± 15.2	86.0 ± 18.0	0.86
SVM - quadratic	87.8 ± 15.0	84.3 ± 16.6	0.86

	IMU-REMY		
Classifier	PPV [%]	Sensitivity [%]	F1-score
RF	74.6 ± 19.1	64.0 ± 27.3	0.69
kNN	75.1 ± 19.8	69.8 ± 24.5	0.72
SVM - linear	56.1 ± 27.3	35.1 ± 29.1	0.43
SVM - gaussian	81.2 ± 18.7	58.4 ± 26.4	0.68
SVM - quadratic	82.7 ± 15.3	65.9 ± 28.7	0.73

5.3 Results

Table 5.3 shows the results of the LOPO validation on the IMU-LAB and IMU-REMY datasets. Shown is the average performance of the classifier and the standard deviation. For the kNN classifier, a value of $k = 5$ resulted in the highest performance. For all classifiers, the window length was varied between 5-20s. The highest combination of classifier scores was obtained with a window size of 15 s. For the IMU-LAB datasets, all classifiers perform similarly. On the IMU-REMY data, the SVM with quadratic kernel has the highest F_1 score of 0.73, resulting in an average PPV and sensitivity of 82.7 % and 65.9 % respectively.

Figure 5.8 shows the effect of the window length on the average PPV and sensitivity of the quadratic kernel SVM classifier validated on the IMU-REMY dataset. While a maximum PPV presents itself at a window length of 19 s, a window length of 15 s results in the highest combination of PPV and sensitivity, with only a marginal decrease in PPV.

Figure 5.9 shows the PPV-sensitivity curves of all classifiers on the IMU-REMY dataset. These were obtained by concatenating the scores of all iterations of the LOPO validation. The gaps visible in the kNN and RF classifier are the result of no scores existing in those areas.

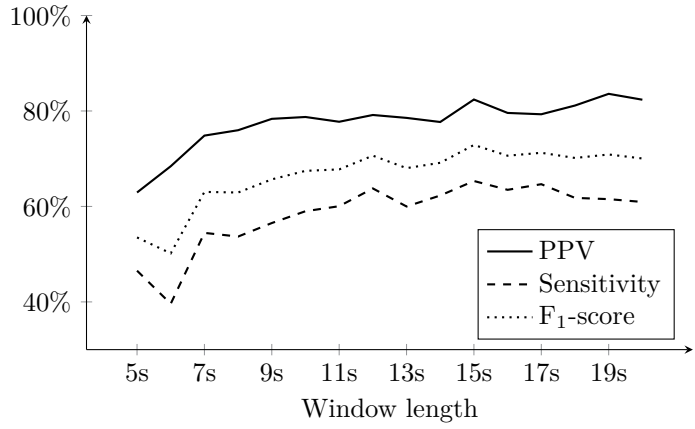


Figure 5.8: Effect of the window length on performance for the SVM classifier with quadratic kernel on the IMU-REMY dataset.

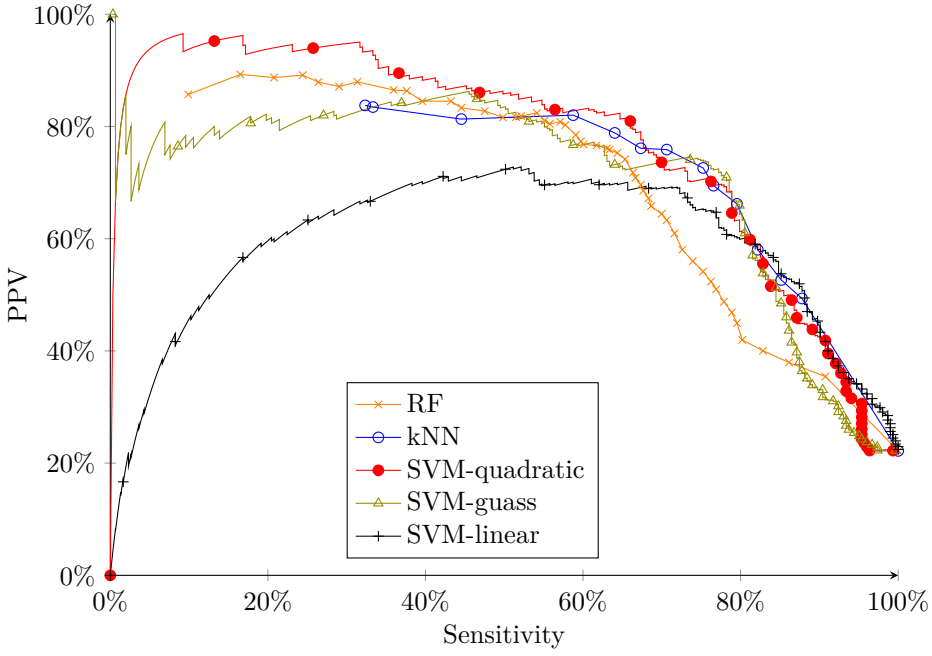


Figure 5.9: PPV-sensitivity curves for each of the evaluated classifiers validated using the IMU-REMY dataset.

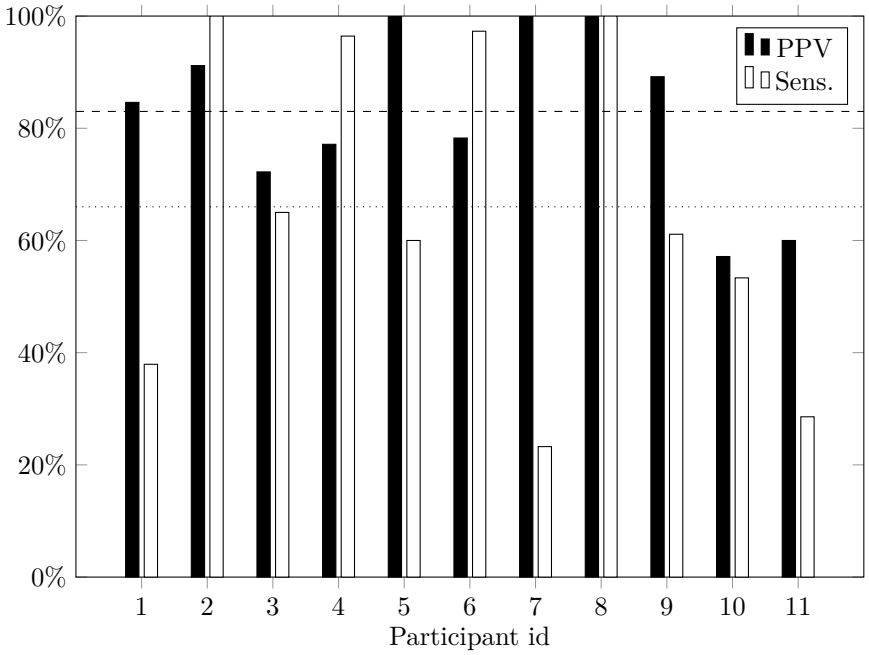


Figure 5.10: PPV and sensitivity of the quadratic SVM for each participant in the IMU-REMY dataset. The dashed and dotted horizontal lines denote mean PPV and sensitivity over all participants, respectively.

Figure 5.10 shows a bar graph with the PPV and sensitivity scores of the quadratic SVM for each participant individually in the IMU-REMY dataset.

5.4 Discussion

The kNN and quadratic kernel SVM classifier have almost the same F1-scores and both perform well on the IMU-REMY dataset, each achieving the highest sensitivity and PPV respectively. When choosing the best classifier, a trade-off between PPV and sensitivity will have to be made. The highest F1 score of 0.73, which is the highest equally weighed combination of PPV and sensitivity, is achieved by the SVM with quadratic kernel. As seen in Figure 5.9, however, a different operating point can result in another classifier achieving a higher performance, depending on the chosen trade-off between PPV and sensitivity.

Around recall values of 78 %, for example, the Gaussian kernel outperforms the quadratic kernel for the real-life dataset, although the difference is marginal.

As seen in Table 5.3, there is little to no variation between classifier types trained on the IMU-LAB dataset. When data is recorded in a controlled environment, epochs of chewing and not-chewing typically don't contain any auxiliary movements that can skew performance of the features and as such, different classifier types perform equally well. As expected, performance drops when transitioning to a real-life environment with the IMU-REMY dataset. Auxiliary movements performed while both chewing and not-chewing will have a negative impact on classifier performance. Since subjects were often recorded in pairs of two, talking is one of the auxiliary activities predominately present in the data.

A maximum performance with a window length of 15 s is obtained using the presented approach, as shown in Figure 5.8. This length corresponds to approximately half the average duration of chewing epochs in the IMU-REMY dataset. The average time spent chewing a bite is relatively high. This is because several participants in the IMU-REMY dataset had the tendency of taking little to no pauses between bites and eat the meal while almost continuously chewing.

Figure 5.10 shows that there is high variability on performance between participants in the IMU-REMY dataset. Variability is smaller in the IMU-LAB dataset, as shown by the smaller standard deviation on the mean performance metrics shown in Table 5.3. This is to be expected, as the controlled environment of the IMU-LAB dataset makes measurements among participants more uniform. In the IMU-REMY dataset, there are some participants with perfect performance scores. Participant 8 has a perfect classification performance of 100 % PPV and sensitivity. Participants 5 and 7 have 100 % PPV, with participant 2 furthermore having 100 % recall. Six participants are above the mean PPV value over all participants, while only 4 are above the mean sensitivity value.

Participants 1, 7 and 11 show sensitivity scores below 40 %. A low sensitivity means that the algorithm missed many segments of chewing and incorrectly labelled them as not-chewing. There are two reasons for this. Participants 1 and 11 were very active during the meal, with lots of talking with their neighbour during the meal and movements of the head. Because of this, the auxiliary movements are dominant in the acceleration signal and the algorithm has trouble to detect the chewing motion during these periods. Furthermore, the low performance of participant 7 is believed to be caused by the type of food in combination with auxiliary movements. This participant was served a meal with very soft foods and did little chewing during the meal, mostly directly swallowing the food, and also showed lots of head movement during the meal.

The system presented by Sazonov and Fontana [18] with jaw strain gauge is closest related to ours. It also uses a wearable device for the detection of chewing motion by measuring movements or vibrations of the head during mastication. They do not report PPV and sensitivity scores, but calculated the accuracy of their system as $(PPV + sensitivity)/2$. They report an accuracy of 81 %. If we do the same for the method presented in this thesis, we obtain a maximum accuracy of 74.5 % on the IMU-REMY dataset. While their system outperforms ours in terms of overall detection accuracy, our system utilises only a single accelerometer that can easily be integrated into the elderly's routine and doesn't require any sensors taped to the body. The audio based system demonstrated by Passler *et al.* [35] uses a microphone integrated in a hearing aid capable of detecting chewing activity with a PPV and sensitivity of 91.3 % and 81.8 % respectively. Their work, however, is difficult to compare to ours, as they also include data from young adults and data recording took place in a controlled environment where participants were not allowed to talk during meals. If we compare the results of our IMU-LAB dataset to theirs, ours outperforms the audio system by 3 % in terms of sensitivity and is behind by 4 % on PPV. Lopez-Meyer *et al.* [33] follow a similar approach using a throat microphone. They do not report PPV and sensitivity scores, but are able to achieve a detection accuracy of 94 %. Their study, however, did not include older adults and rely on an obtrusive throat microphone.

5.5 Strengths and Limitations

The aforementioned studies all target healthy young to middle-aged adults and rely on stigmatising or uncomfortable wearables. The system presented in this thesis is more suited for use by older adults. There is, however, still room for improvement. One limitation of the system is the difficulty in distinguishing chewing activity from activities with lots of head movements. Because the vibrations and movements caused by chewing have a low magnitude in the acceleration signal, they can be 'overpowered' by other activities.

Furthermore, our feature set is relatively small and the majority of features are derived from the time domain signal. Performance may be improved by selecting more or better features. Distinguishing between chewing and other activities can be improved by utilising more features from the frequency domain. These features may be more immune to movements of the head. A larger dataset can also result in a better performing generalised model. Furthermore, the influence of body movements on the detection accuracy may be reduced by supplementing the system with a wrist mounted IMU. This would of course

increase complexity of the system, while older adults prefer the system to be as simple as possible.

Another limitation that was encountered during the data collection process is that of personalised meals. In home care or nursing homes, older adults are typically already served a meal that is somewhat personalised. In the event of a decreased bite strength or wearing of dentures, meals may be supplemented or replaced entirely with soft of mixed foods [95]. One participant in our study was served such a meal consisting of partly-mixed foods and as such did not have to chew much. When this is the case, the system has difficulty to detect the food intake because little to none chewing can be detected.

The sensitivity of the current system may not be suitable for use in clinical practice, where a high sensitivity is typically required. However, a high sensitivity was obtained in a controlled environment. This can indicate that the system has the best application potential in research environments, where it is less of an issue to place some limitations on the test subject.

User acceptance of the prototype capture system in its current form was high. None of the participants had a problem wearing a sensor on their glasses, nor did anyone experience discomfort due to the light nature of the sensor. The majority of participants would prefer a more integrated and invisible system. Possible form factors include a clip-on system or a type of sleeve that slides over the frame, which would be virtually invisible and light enough not to be noticeable by the wearer.

5.6 Chewing Strength

The primary goal of the work was to test the feasibility of the accelerometer to measure chewing activity. The system, however, could be used for additional evaluation of the chewing activity. For example, difficulty with chewing has been shown to be prevalent among older adults due to a decreased muscle strength or the use of dentures [96–98]. Using electromyography (EMG) recordings with sensors attached to the temporalis and masseter muscles, a decreased chewing strength in older adults compared to younger adults can be measured [99]. Our proposed accelerometer system indirectly captures some parts of the movements from these muscles during chewing. The accelerometer signal could thus also contain information about the chewing strength of the subject.

Figure 5.11 shows a segment of an accelerometer signal from the IMU-LAB dataset. Shown is a segment of data where the test subject consumes a single bite from a sandwich. The chewing activity is visible in the signal in the form

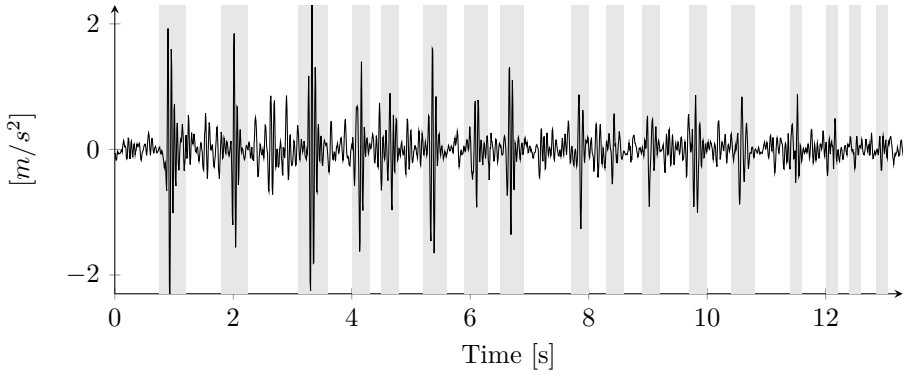


Figure 5.11: Illustration of an accelerometer signal captured while chewing in a controlled environment (data from IMU-LAB). Marked in grey are the peaks that correspond to the closing of the jaw when the subject bites down on the food. Only the x-axis of the IMU is shown.

of high amplitude peaks. The shaded areas indicate these kind of peaks that correspond to the closing of the jaw as the subject bites down on the food. This was validated using the synchronised video recording. The amplitude of the peaks shows a downward trend towards the end of the segment. This may indicate that a correlation exists between the amplitude of the signal and chewing strength. Food becomes softer as the subject chews, lowering the amount of force required to chew. The highest amplitude peaks can be observed at the start of the segment, when the piece of food is still in its hardest form, requiring more strength to break up the food.

This kind of trend is observed regularly in the IMU-LAB dataset, which was captured in a controlled environment, but is no longer visible in the IMU-REMY dataset. This can be due to the fact that the older test subjects in IMU-REMY already had a lower chewing strength, thus making the chewing activity less pronounced. Another possible cause is the uncontrolled nature of the dataset. As users interact with the environment while chewing, the chewing motion is no longer clearly visible in the time domain, as discussed in Sections 5.4 and 5.5.

Confirming the correlation between signal amplitude and chewing strength could be the topic of future work. This may be useful in the study of eating behaviour or early detection of malnutrition risk in older adults.

5.7 Conclusion

An algorithm for the detection of chewing motion from a glasses mounted accelerometer has been presented in this chapter. We showed that an accelerometer can be used for this task with high performance, on par with methods presented in the literature, both in a lab and real-life environment. Using an accelerometer can offer benefits over other methodologies presented in the literature. While not all people wear their glasses continuously, our system does provide an alternative to other, more stigmatising, capture methods.

Chapter 6

Design of a Smart Plate

The design methodology of the custom smart plate system that was developed during the PhD research is presented in this chapter. The system consists of an off-the-shelf polymer plate placed on top of a base station. It has the novel feature of being able to measure location of individual bites without embedding sensors into the plate itself, all hardware is contained in the base. The background of the location detection is introduced. Furthermore, the mechanical and electrical design of the system is discussed.

Parts of the work presented in this chapter have been published in the following paper:

- Gert Mertes, Glenn Christiaensen, Hans Hallez, Sammy Verslype, Wei Chen, and Bart Vanrumste. Measuring weight and location of individual bites using a sensor augmented smart plate. In *Proc. of the 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2018.

6.1 Introduction

Table embedded scales such as the Universal Eating Monitor (UEM) have been introduced in the literature for the measurement of the amount of food consumed during meals. Embedding sensors into the eating surface allows the user to simply consume the meal without manually having to measure the weight of food before and after the meal. A processing system can then automatically

extract the beginning and end weights of the meal. While this may seem like a small benefit, recall that requiring manual data entry is one of the major causes of users not adhering to food intake monitoring protocols [10, 31].

To date, several research groups have been active in the modernisation of the UEM concept and a number of different adaptations to the UEM have been presented in recent literature [27–30, 45, 46]. All of these systems, however, only measure the total weight profile of the meal and do not have the ability to measure the type of food that is being consumed. One of the goals of food intake monitoring is to measure the total amount of ingested calories, which can differ greatly between food types [100]. The amount of ingested calories can be estimated from the total weight by comparing it to the total calorie content of the entire meal, which may be known in some cases, for example when food is being served from a central kitchen. Being able to detect individual food types would allow for a more detailed calorie measurement. The measured consumed weight of each food type can be compared to a database of known foods, containing calories per unit of weight, to estimate the caloric content of the meal.

Zhou *et al.* [30] propose a weighing surface with a force sensitive pressure matrix embedded in the table cloth. This allows for the detection of different eating implements on the table, such as the location of the plate or drinking cups. This system can detect when implements are removed and can, for example, be used to detect fluid intake. The detected removal of a cup can be combined with the measured decrease in weight to estimate fluid intake, but individual food types remain undetectable.

A potential method for measuring individual food types, is by serving each food type in a different container and measuring the decrease in weight from each container during the meal. This still requires some manual input by the user, as food types need to be assigned to each container, but it can be an improvement over existing manual measurement techniques. Chang *et. al* [28] propose a table with multiple independent weighing surfaces. Food can be served on these weighing surfaces, and the decrease in weight can be detected from each surface individually. This requires each food type to be served on an individual plate or bowl and requires extensive adaptation of the eating surface.

Given a compartmentalised plate, such as those found in home food delivery services or hospitals, food types may be served in an individual compartment. By detecting weight changes in each compartment, individual food types can be measured, given that there is prior knowledge of which food is in each compartment and that compartments are not mixed together. Weight sensors may be embedded in each compartment of the plate to perform the measurements. This approach would require the development of a custom plate,

with sensors and a processing system embedded into the plate itself. While this is technically possible, it introduces a number of challenges. First and foremost, it would require a robust electrical design that can, at the very least, withstand high temperatures, is waterproof or dishwasher resistant and may even need to be microwave resistant. Furthermore, users would need to buy multiple, possibly expensive, systems if they wanted more than one plate.

A different approach could be through a base station with different weighing surfaces, similar to the dining table proposed by Chang *et. al* [28], but in the form of a single plate. Small plates with the size of each weighing surface could then be used to serve individual food types and placed on top of each weighing surface on the base. While this simplifies the hardware requirements as the plates themselves don't need any sensors built in, it would still require the manufacturing of highly customised plates to be used only with this system.

We propose a novel approach to this problem that works with an off-the-shelf plate with 3 compartments. All hardware is contained in a base station, but the base does not contain different weighing surfaces. Instead, the plate is placed on top of multiple weight sensors and by looking at the changes in weight distribution among all sensors during eating, the location of bites can be detected. By comparing the position of each bite to the layout of the plate, compartments can be assigned to each bite. This approach does not require a customised plate or eating surface.

The mechanical and electrical design and implementation of the smart plate was performed in collaboration with master thesis student Glenn Christiaensen.

6.2 Bite Localisation

The bite localisation algorithm is based on changes in the Centre of Pressure (COP) of the plate during eating. Bites are located on the plate in the horizontal plane. Each bite is assigned an x and y coordinate that denotes its position on the plate relative to a chosen origin. To measure COP in two dimensions, at least three weight sensors are required. This section introduces the formulas used for the localisation. To simplify the deduction, a system with two weight sensors and localisation in one dimension is shown, but this is later easily extended to two dimensions.

Figure 6.1 depicts a static force system with a mass m_1 at unknown position x_1. Assume that force sensors are embedded into points A and B, making the reaction forces F_A and F_B known, and the horizontal beam is the weighing

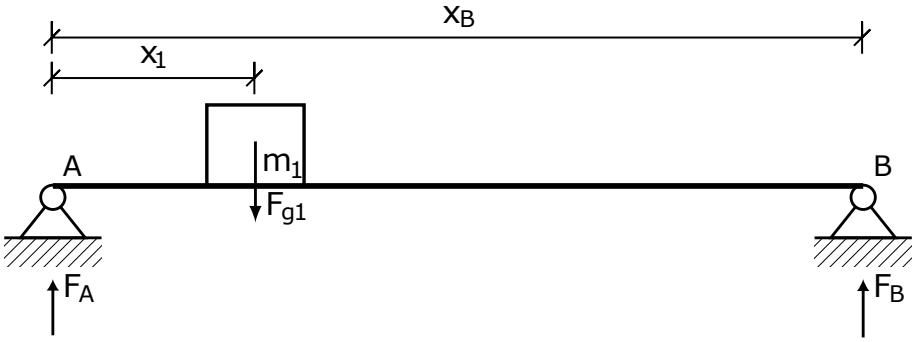


Figure 6.1: Static force system with 1 mass.

surface. The location x_1 of the mass can be found by substituting the force equilibrium into the momentum equilibrium formula around point A:

$$\begin{cases} F_{g1} & = -F_A - F_B \\ x_1 F_{g1} & = -x_B F_B \end{cases} \rightarrow x_1 = \frac{x_B F_B}{F_A + F_B} \tag{6.1}$$

The forces in point A and B are measured with a force sensor and are thus known, as is the distance between both sensors x_B. The formula for x_1 is thus fully defined and corresponds to the COP of the static system.

We now introduce a second mass m_2 with a different and unknown location x_2. This is shown in Figure 6.2. The location x_2 cannot be determined using the classic COP formula, because the COP will be somewhere in between both masses. Similar to x_1, the location of x_2 can be derived through substitution of the force equilibrium and momentum equilibrium formula around point A:

$$\begin{cases} F_{g1} + F_{g2} & = -F_A - F_B \\ x_1 F_{g1} + x_2 F_{g2} & = -x_B F_B \end{cases} \rightarrow x_2 = \frac{x_B F_B + x_1 F_{g1}}{F_A + F_B + F_{g1}} \tag{6.2}$$

This formula for x_2 is not fully defined, because F_{g1} and x_1 cannot be derived from the *current* values of F_A and F_B. However, the values of F_A and F_B *prior* to the introduction of m_2 can be used to determine F_{g1} and x_1.

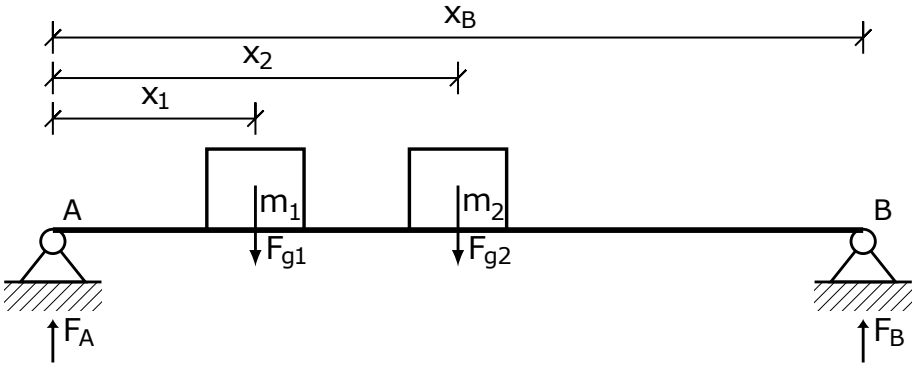


Figure 6.2: Static force system with 2 masses.

Formula 6.2 can be rewritten and simplified to the following form:

$$\begin{aligned} x_2 &= \frac{x_B F_{Bt_1} + x_1(-F_{At_0} - F_{Bt_0})}{F_{At_1} + F_{Bt_1} - F_{At_0} - F_{Bt_0}} \\ \Leftrightarrow x_2 &= \frac{x_B F_{Bt_1} + \frac{x_B F_{Bt_0}(-F_{At_0} - F_{Bt_0})}{F_{At_0} + F_{Bt_0}}}{F_{At_1} + F_{Bt_1} - F_{At_0} - F_{Bt_0}} \\ \Leftrightarrow x_2 &= \frac{x_B F_{Bt_1} - x_B F_{Bt_0}}{F_{At_1} + F_{Bt_1} - F_{At_0} - F_{Bt_0}} \end{aligned} \tag{6.3}$$

with t_0 and t_1 indicating the points in time before and after the introduction of m_2, respectively.

By substituting $(t_1 - t_0)$ with Δt, x_2 is defined as:

$$x_2 = \frac{x_B \Delta F_B}{\Delta F_A + \Delta F_B} \tag{6.4}$$

with ΔF the difference in force before and after introduction of m_2.

The location of a second mass can thus be found using the difference in weight on points A and B before and after introduction of the mass. To link back to our smart plate: the mass m_1 is food that is already present on the plate and m_2 is a single bite removed from the plate, in the form of a negative mass added to the system.

Finally, the derivation can also be done for a two-dimensional system and formula 6.4 can be rewritten to allow for multiple weight sensors, resulting in:

$$\begin{cases} x = \frac{1}{\Delta F_{\text{tot}}} \sum_{i=1}^{n} \Delta F_i x_i \\ y = \frac{1}{\Delta F_{\text{tot}}} \sum_{i=1}^{n} \Delta F_i y_i \end{cases} \tag{6.5}$$

with x and y the position of the bite relative to a chosen origin, ΔF_{tot} the total change in force on all force sensors before and after the bite, n the total number of force sensors, ΔF_i the force difference measured by sensor i and x_i and y_i the distance from force sensor i to the chosen origin.

To localise individual bites, a bite detection algorithm is required to first locate each bite in time, calculate the change in force before and after the bite and enter that into Formula 6.5. This bite detection algorithm is introduced in Chapter 7, together with the validation of the system on real-life data.

6.3 Strain Gauge Load Cell

Load cells are a commonly used sensor when an accurate measurement of force or weight is required [101]. Load cells generate an electrical signal with a magnitude that is directly proportional to the applied force. After calibration, they can thus be used to measure weight. There are several different types of load cell, but for small scale weighing applications, the strain gauge load cell is commonly used due to a low cost and high sensitivity.

The strain gauge load cell works by measuring the deformation of a structural member as a result of an applied load. Figure 6.3 shows a drawing of a beam-type load cell. The beam is made of a stiff material, aluminium is commonly used, that is fixed to a base on one end. When a load is applied on the free end, the beam will deform slightly. Strain gauges are attached at different points to the beam. As the beam deforms, the strain gauges will experience tension or compression, depending on their location. In the configuration depicted in Figure 6.3, two cells will experience compression while two will be under tension. When strain gauges deform, their electrical resistance changes, by an amount that is proportional to the deformation. The electrical resistance of the strain gauges is thus a measure for the physical deformation of the beam, which in turn is a measure for the applied load.

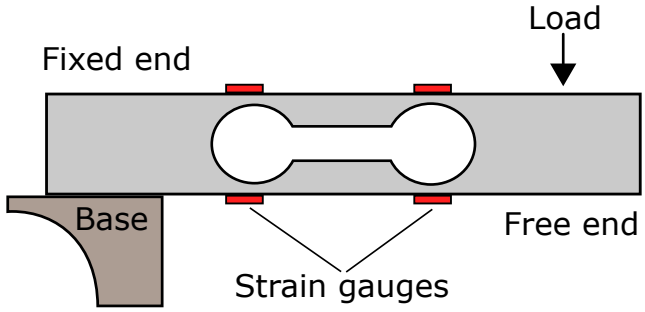


Figure 6.3: Drawing of a strain gauge load cell.

By placing the strain gauges in a Wheatstone bridge, the change in resistance can be converted into a voltage. As two strain gauges compress, their resistance decreases, causing more current to flow in them. The other two strain gauges are stretched, which increases their resistance, decreasing their current flow. A voltage difference can thus be measured at the output that is proportional to the applied load. Using four strain gauges results in a high sensitivity and has the added benefit that the load cell is temperature compensated. Each strain gauge will expand or compress equally under temperature changes and the differential signal will remain the same.

Strain gauge load cells of the beam-type are used in the design of our smart plate. The CZL616C load cell is used because of its small size and high accuracy. This is a readily available load cell with a rated capacity of 780 g and is accurate to <1 g, which is more than sufficient for our application. In our prototype, the weight of the meal is distributed over three load cells, allowing for a total rated capacity of 2340 g. The important characteristics of the load cell are shown in Table 6.1. These properties are:

Rated Capacity The maximum nominal load the cell is designed for.

Safe Overload The maximum force that can be applied without changing the mechanical and electrical characteristics of the cell.

Drift The maximum amount that the output will drift over time given a constant applied force.

Non-linearity Maximum deviation from the linear calibration curve given a two-point calibration.

Non-repeatability Maximum difference between two measurements with the same load applied.

Table 6.1: Mechanical and electrical characteristics of the CZL616C load cell.

Description	Value
Rated Capacity	780 g
Safe Overload	120 % rated capacity
Drift	1.6 g/hr (max)
Non-linearity	390 mg (max)
Non-repeatability	390 mg (max)
Supply voltage	5.0 V DC (max)
Output	$(800 \pm 100)\,\mu V/V$
Dimensions	45 mm × 9 mm × 6 mm

Supply Voltage The load cell needs to be energised in order to work, this is the maximum voltage that can be applied.

Output The cell will produce this value at the output when loaded at rated capacity for each V of supply voltage.

6.4 Mechanical Design

Figure 6.4 shows a picture of the smart plate system. The load cells, embedded system and battery hardware are located in the base station underneath the polymer plate. The plate has three compartments and resembles those that are typically used in hospitals or nursing homes. The plate is mounted on top of three load cells and is easily attached and detached. In the current prototype, a reversible friction fit connection is used that only works with this particular type of polymer plate, but in the future this could be extended to allow different plate types. The base platform is made from acrylic, which is lightweight and easily cleaned.

Figure 6.5 shows a top view of the load cells as they are mounted to the base platform. The cells are placed in increments of 120° on the circumference of the base. One end of the load cell is permanently fixed to the base platform while the other end is free floating, in a similar configuration as shown in Figure 6.3. The polymer plate is mounted directly on top of the free ends of the load cells. The total weight of the plate is thus transferred to and spread over the three cells. By using three mounting points, the round plate is always stable.



Figure 6.4: The smart plate system with polymer plate and base station.

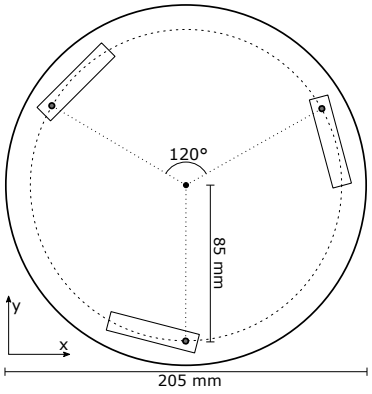


Figure 6.5: Position of the load cells on the base platform.

6.5 Electrical Design

A custom embedded system was designed for the smart plate prototype. While there are some commercially available products that can capture the signals coming from the load cells, none fit our requirements. To be a standalone data collection system, the system should have the capacity to store the data locally on a persistent storage medium, such as an SD card. Commercial products require interfacing with a computer to store the data. The choice was therefore made to design our own embedded system to read and store the load cell signals. This also allowed us to build in features that are not available in commercial products, such as Bluetooth connectivity.

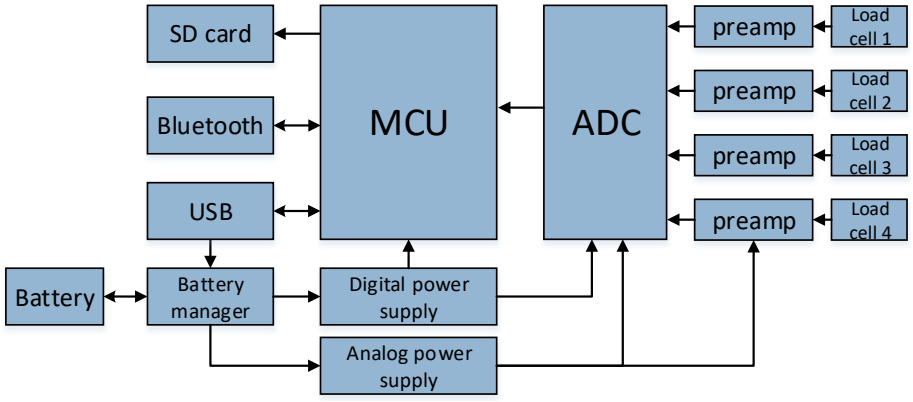


Figure 6.6: Block diagram of the embedded capture system.

6.5.1 Embedded System

Figure 6.6 shows a block diagram of the embedded system. The system is controlled by an 8-bit Atmel ATmega644 microcontroller. A relatively low performance controller was chosen because the system does not need to perform computationally intensive tasks. The data only needs to be stored on a persistent medium and can be relayed to a different system for processing. There are 4 independent load cell channels, each with their own pre-amplifier, that are simultaneously sampled by the Analogue to Digital Converter (ADC). The ADC is a 24-bit 4-channel ADS1224. Note that in the current smart plate prototype, only 3 load cells are used. Data is sampled at a selectable sample rate of 6 Hz, 12 Hz or 24 Hz.

Interfacing with the smart plate can be done via a USB connection or via Bluetooth. In both cases, communication is handled via a serial connection. The Bluetooth module implements the Serial Port Profile (SPP) while an FTDI FT230x allows serial communication via USB. The firmware implements a text based command protocol that allows commands and data to be sent over the serial link. Data from the load cells can be stored on the SD card, streamed via USB/Bluetooth or both simultaneously. The system is powered by a single cell 2000 mAh Lithium-ion battery. The USB port is used to charge the device. While charging, the system can remain operational with power taken from the USB port instead of the battery. Power usage of the device is low. Data can be continuously sampled and stored on the SD card for approximately 48 hours. When used to measure two to three meals per day, the system can be used for approximately 3 weeks without charging.

6.5.2 Power Supply

To achieve maximum sensitivity, the load cell should be powered at its maximum supply voltage of 5.0 V. To minimise power usage, however, the supply of the digital system should be as low as possible. While the MCU and ADC support supply voltages down to 2.7 V, the SD card requires at least 3.3 V. Furthermore, for maximum signal integrity, it is good design practice to keep the analogue and digital supply rails of the system separate, as the clock from the microcontroller may introduce noise in the supply rail.

Two independent power supplies are therefore implemented in the system: a supply of 3.3 V for the digital domain and one of 5.0 V for the analogue domain. Because the LI-on battery voltage can vary between 2.7 V and 4.2 V, the digital power supply is of the buck-boost switch mode variant, implemented by an LTC3440. The analogue power supply is implemented in two stages. First, the battery voltage is boosted to 5.5 V by another LTC3440. This is then regulated down to 5.0 V by an ultralow noise linear regulator of the type LT3042. This two step regulation is used because the boost-converter will introduce switching noise to the supply rail. Because load cells are very sensitive and only produce a voltage deflection in the range of several hundred microvolts, any switching noise in the analogue supply rail will have an impact on the load cell measurements. The linear regulator is therefore used to eliminate the switching noise and generate a clean supply rail.

6.5.3 Load Cell Preamp

While load cells are passive sensors, they need to be powered with a supply voltage for operation. As discussed in Section 6.3, the cell contains strain gauges that change resistance as they compress and stretch. At the output, the load cell generates a voltage difference that is proportional to the supply voltage. The load cells in our system generate an output of $(800 \pm 100)\,\mu V/V$ when loaded at rated capacity. For a supply voltage of 5.0 V, this results in a maximum output of (4.0 ± 0.5) mV. Because this is the output at rated capacity, the cell will produce less during normal operation, typically around 1 mV.

This output is too small to be measured by the ADC, so amplification is required. Each load cell has its own independent preamp. Amplification is done by a low-noise AD623 differential instrumentation amplifier with a gain of 1000. This chip is commonly found in precise load cell applications, such as high-resolution weight scales. Figure 6.7 shows the preamp of a single load cell. The cell is represented by resistors (the strain gauges) placed in a Wheatstone bridge. AVDD is the analogue supply rail. While the differential amplification filters

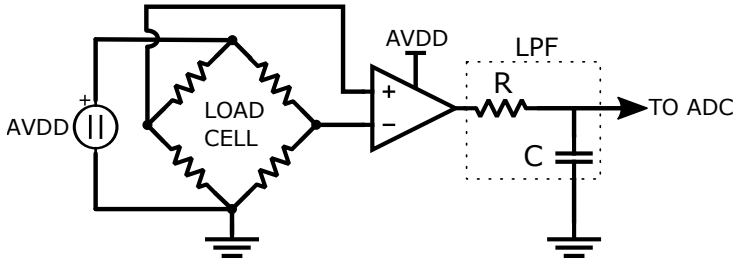


Figure 6.7: Differential amplifier and low pass filter for the load cell.

out any noise picked up by the load cell, a first-order Low Pass Filter (LPF) with a cut-off frequency of 1 kHz is placed after amplification for good measure. The filtered signal is then sent to the ADC.

6.5.4 PCB

Due to the noise sensitive nature of the load cells that generate only a few mV during normal operations, care was taken during the Printed Circuit Board (PCB) design process to maintain signal integrity. A four-layer PCB was designed. Signal layers are located on the top and bottom layer, while a power and ground plane is placed in the middle. This approach minimises capacitive coupling and prevents ground loops from forming. Furthermore, the digital and analogue domain is physically separated on the PCB to further reduce crosstalk and maintain signal integrity in the analogue domain. The power and ground planes are also physically split up so that each domain has its own dedicated plane. Figure 6.8 shows a picture of the final PCB.

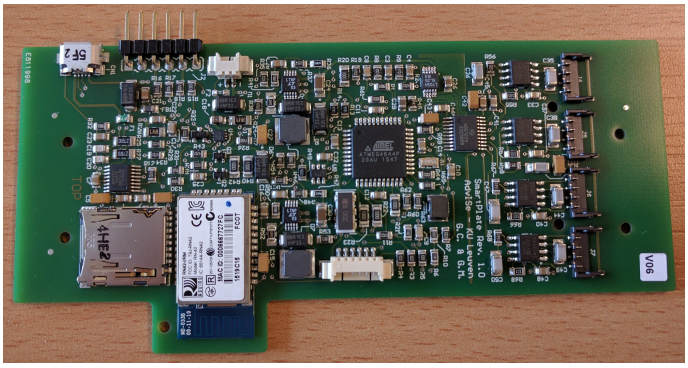


Figure 6.8: The custom PCB design for the smart plate.

6.6 System Tests

To ensure correct operation of the smart plate system, several system tests were performed. We tested the linearity, repeatability, drift and location accuracy.

6.6.1 Calibration

Prior to usage, the load cells need to be calibrated. This can be done with a two-point calibration, given that the sensors are linear. A calibration mass of 500 g is used, which is the maximum typical load expected during normal operation. Each load cell is calibrated individually by placing the mass directly on the cell. The calibration factor G is defined by:

$$G = \frac{\text{calibration mass}}{ADC_{load} - ADC_{zero}} \tag{6.6}$$

with ADC_{load} the output of the ADC when the calibration mass is applied and ADC_{zero} the output without any load applied. The weight m can then be calculated from any given ADC value with:

$$m = G.(ADC - ADC_{zero}) \tag{6.7}$$

6.6.2 Linearity

After two-point calibration, the linearity of the load cells is tested by placing a known mass in the middle of the plate and measuring the output. The results are shown in Figure 6.9. Figure 6.9a shows a plot of the measured weight against the true weight. The measured weights follow the reference line perfectly, indicating linearity. Figure 6.9b shows the error in grams per measurement. A single outlier of 0.5 g is observed at 1500 g, but overall the error is negligible. The test shows that the system is linear well within the advertised specifications of the load cells and that a two-point calibration is sufficient.

6.6.3 Drift

When a load is applied for long periods of time, the output of a load cell tends to drift. This was tested by placing a mass of 200 g in the middle of plate and sample the weight for 8 hours. The result is shown in Figure 6.10. After 8

hours, the weight has drifted upwards by 2.5 g. This corresponds to a drift of 0.31 g/h, which is well under the maximum advertised drift of the load cell and negligible given that typical meals last less than one hour and a single bite is typically around 5 g to 10 g [29, 102].

6.6.4 Repeatability

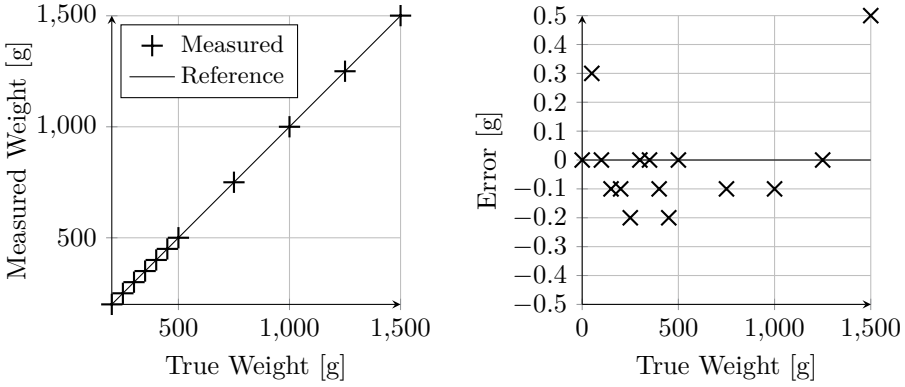
To test the repeatability of the system, a mass of 200 g is repeatedly placed on the plate. After placing the mass, its weight was measured by the system, and the mass was removed again, allowing the system to stabilise before the next measurement. Figure 6.11 shows that the system is repeatable well within the advertised load cell characteristics.

6.6.5 Position Accuracy

To evaluate the accuracy of the position detection, a mass of 5 g is placed at predefined positions on the plate. The location of the mass is then measured using the formula introduced in Section 6.2. Recall that the location detection needs the measured weights on all load cells before and after the mass is placed (or removed). For this test, the load cells are continuously sampled while the mass is placed at different locations. Afterwards, the weights before and after each introduction of the mass are manually entered into the formula. The results are shown in Figure 6.12. This shows a top-down view of the detected position on the plate. The location where grid lines intersect denote the actual position of the mass, while the crosses indicate measured positions. The black dots indicate the position where the plate is attached to the load cell. The system is able to measure the location of the mass with a mean error of around 5 mm. For the purpose of assigning detected bites to compartments, this is accurate enough.

6.7 Conclusion

The hardware of the custom smart plate system has been introduced in this chapter. The background of the bite localisation has been shown together with the design considerations of the mechanical and electrical design. Tests performed with the system show that it performs as expected and well within the advertised characteristics of the load cells.



(a) Linearity (b) Weight error

Figure 6.9: Linearity test.

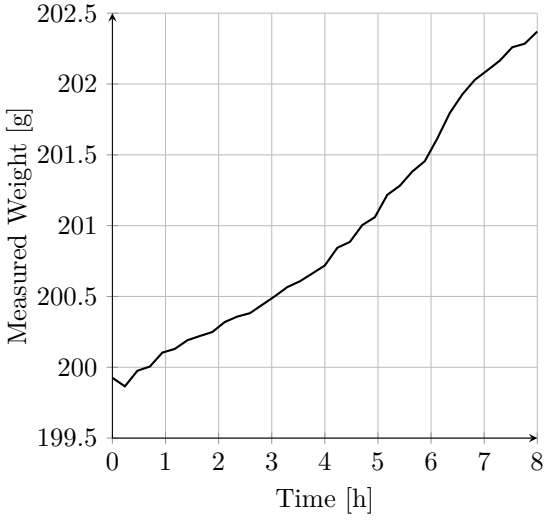


Figure 6.10: Drift of the system during 8 hours of constant load.

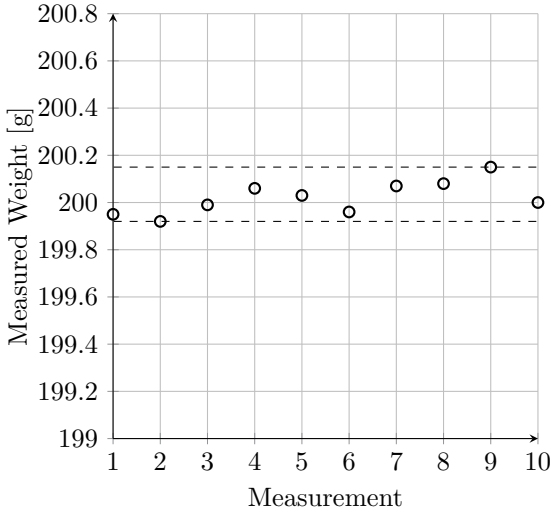


Figure 6.11: Repeatability of the system with a weight of 200g. Total deflection is 0.23 g.

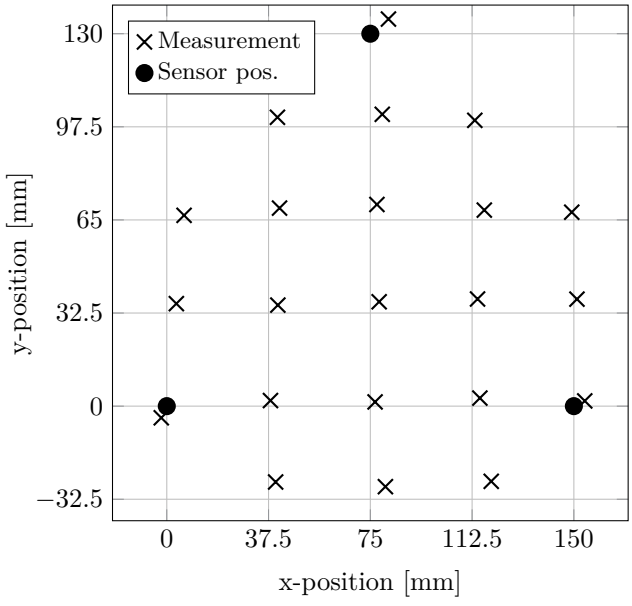


Figure 6.12: Position test with a weight of 5 g. The locations where grid lines intersect indicate the true locations of the weight.

Chapter 7

Detection of Food Bites Using a Smart Plate During Unrestricted Eating

The previous chapter introduced the smart plate system and the bite localisation methodology. In this chapter, the bite detection algorithm required for the bite localisation is presented. An algorithm based on a supervised classifier is proposed. The challenges related to measuring individual bites during unrestricted eating are presented. The system is tested on data recorded from older adults in a nursing home and hospital. The results of both the bite detection and bite localisation are discussed.

The preliminary study discussed in Section 7.2 of this chapter has been published in the following paper:

- Gert Mertes, Glenn Christiaensen, Hans Hallez, Sammy Verslype, Wei Chen, and Bart Vanrumste. Measuring weight and location of individual bites using a sensor augmented smart plate. In *Annual Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*. IEEE, 2018.

Furthermore, the new bite detection algorithm has been submitted and is at the time of writing under review in the following paper:

- Gert Mertes, Li Ding, Wei Chen, Hans Hallez, Jie Jia and Bart Vanrumste Measuring and localizing individual bites using a sensor augmented plate during unrestricted eating for the aging population *IEEE Journal of Biomedical and Health Informatics.* (in revision).

7.1 Introduction

In its most basic modus of operation, the Universal Eating Monitor (UEM) can be used to automatically measure the beginning and end weights of a meal to measure the total amount of consumed food. These measurements can then be used to track food intake over time, without requiring any input from the user. In situations where manpower is expensive, such as a hospital, this kind of system can already provide benefits over existing, more time-consuming, manual methods. Furthermore, automating these kind of measurements can increase user adherence. This information can also be measured with the smart plate system.

While this relatively simple measurement of consumed food can already provide benefits over traditional monitoring methods, additional parameters may be extracted from the smart plate to further improve existing methods. Because the output of the smart plate shows the food intake over time, it can be used to calculated cumulative intake curves. These can be used beyond regular food intake monitoring and can provide researchers or clinicians with more detailed information on the effects of psychological, nutritional or pharmacological changes on human appetite [46–48].

In addition to a cumulative intake curve that provides information over the course of the entire meal, even more information can be extracted by detecting the food intake down to the level of each individual bite. Detecting individual bites during the meal provides in depth information about the microstructure of ingestive behaviour, such as average bite size, length, time between bites and number of total bites [49,50]. This information can further be used to study the effects of disease or psychological changes on intake behaviours. The algorithm proposed by Mattfeld *et al.* [29] can detect individual bites from the output of the UEM. Bites can be located in time and their individual weight can be measured. The bite detection algorithm works by comparing the changes in weight to a set of predefined rules. These rules were manually created based on observations in the data. The limitation of this approach is that rules have to be manually defined and may not generalise all bites in the data.

Accurately measuring individual bites is a requirement for the correct operation of the bite localisation algorithm with our custom smart plate. Recall that in order to localise bites on the plate, the force on all loads cells before and after the bite has to be known, as outlined in Section 6.2. In this chapter, we propose a novel bite detection algorithm based on a supervised machine learning model that automatically learns the generalised concepts from the training data, as opposed to using manually created rules.

7.2 Preliminary Study

A preliminary bite detection study was performed under my supervision during the master thesis research of Glenn Christiaensen. In his master thesis, the mechanical and electrical design of the smart plate was implemented. To test the feasibility and operation of the system, a preliminary bite detection algorithm was developed and tested on two test meals.

A rudimentary bite detection algorithm was implemented to test the overall functionality of the plate. The assumption is made that every decrease in weight measured by the plate corresponds to a bite. Bites can thus be detected by measuring the *steps* in the signal. The total weight signal (sum of the three load cells) is first filtered with a moving-median filter. This is done to discard any activity prior to the bite that may be caused by pressing down on the plate or movements such as scooping or cutting. Steps in the signal are then detected with a simple edge detector using the local maxima of the derivative of the signal. The size of a step is subsequently calculated by subtracting the mean value of all samples after the step (up until the next step) from the mean value of all samples before the step (up until the previous step). The bite localisation algorithm can then be applied to each step.

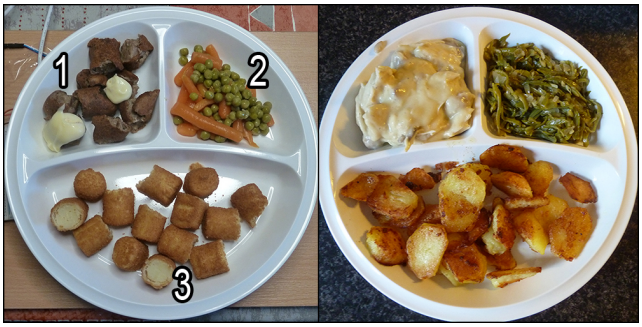


Figure 7.1: The test setup of the preliminary bite detection study. Served food in the lab (left) and home (right) environment is shown. The smart plate compartments are labelled 1 to 3.

The preliminary bite detection algorithm was evaluated on two test meals. Figure 7.1 shows both situations prior to the meal. One subject was asked to consume a meal in a lab environment and an older adult was asked to consume a meal in a home environment. The older adult was recruited from family of the master thesis student. Food was pre-cut in both cases and subjects were only allowed to use a fork. For the lab measurement, the subject was asked to eat slowly and take care to minimise actions that would introduce large

amounts of movement or pressure, such as overly pushing into the plate. For the home measurement, the subject could eat at their own speed and no restrictions were placed on how the food was picked up. Subjects and the plate were video recorded synchronously to the smart plate to validate the detected bites. This was done by synchronising the plate's real time clock with the video capture system clock at the start of each measurement.

The results of both measurements are shown in Table 7.1. Two types of performance metrics can be extracted from the algorithm. First, the bite detection and localisation are evaluated. This is done by comparing the location of the real bites to the detected bites. Recall that the meals were video-recorded, so the location of real bites on the plate is known. For the lab measurement, only a single bite was misplaced out of a total of 34. A bite from compartment 2 was instead detected in compartment 1, resulting in an error of 1 in both compartments. In the home measurement, a bite from compartment 3 was detected in compartment 1. Furthermore, the bite detection algorithm aggregated two bites together that followed in close succession, thus missing one bite from compartment 3.

The second performance metric is the accuracy of the weight detection. The goal of the system is to detect the location of bites and assign them to a compartment, so that the total weight of consumed food from each compartment can be measured. This is done by summing the detected weights of bites for each compartment. The total weight of the meal is then the sum of the three compartments. The detected weight is compared to the true weight of the meal. The true weight is obtained by weighing the foods in each compartment before and after the meal with a kitchen scale. As seen in Table 7.1, the system can accurately measure weights per compartment, with a maximum absolute error of 6.1 g for our test setups. For typical meals in the range of 200 g-400 g, the error is negligible. The error also depends on what kind of error was made during the bite localisation phase. For the lab measurement, a single bite was misplaced, but the weight of that bite was still measured correctly, resulting in a low error on the total measured weight. Similarly, in the home measurement two bites were aggregated together in compartment 3, but their combined weight was still correctly measured. This resulted in a missed detection for that compartment, but a correct measured weight.

The preliminary study shows that the bite localisation and smart plate system operate as expected. Because this was a study in a controlled environment, however, one very important assumption was made: that *every* decrease in weight corresponds to a bite. In practice, this is absolutely not the case. Food can be picked up and replaced on the plate, utensils can be rested on the plate, or the user themselves may lean on the plate. A more advanced bite detection algorithm is thus required for real-world usage of the smart plate.

Table 7.1: The results of the preliminary bite detection tests.

(a) Lab data

Bite Detection

C[a]	Real Bites	Detected Bites	# Errors
C1	9	10	1
C2	11	10	1
C3	14	14	0

Weight Detection

C[a]	Real Weight	Detected Weight	Abs. Error
C1	77.1 g	80.5 g	3.4 g
C2	62.9 g	56.8 g	6.1 g
C3	137.2 g	139.8 g	2.6 g
Tot.	277.2 g	277.1 g	0.1 g

(b) Home data

Bite Detection

C[a]	Real Bites	Detected Bites	# Errors
C1	8	9	1
C2	8	8	0
C3	11	9	2

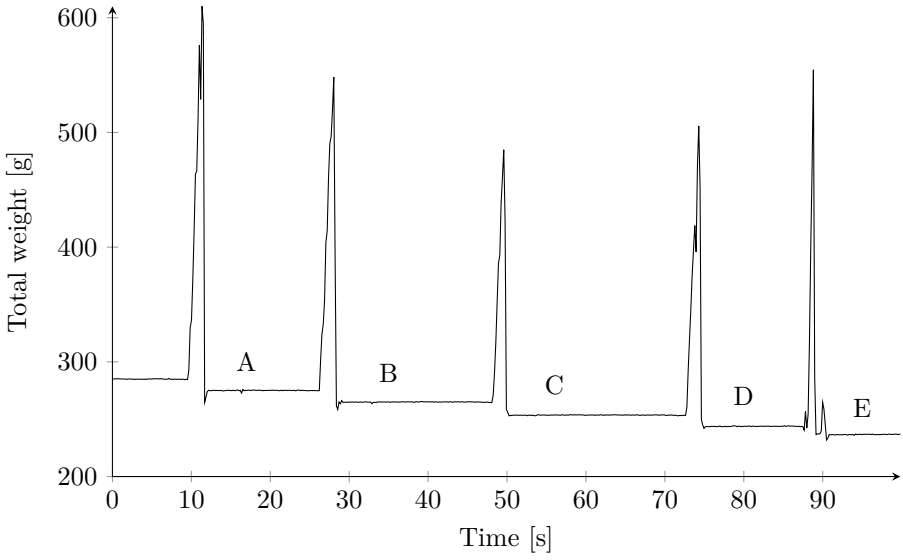
Weight Detection

C[a]	Real Weight	Detected Weight	Abs. Error
C1	115.0 g	114.6 g	0.4 g
C2	102.4 g	100.9 g	1.5 g
C3	64.9 g	65.0 g	0.1 g
Tot.	282.3 g	280.5 g	1.8 g

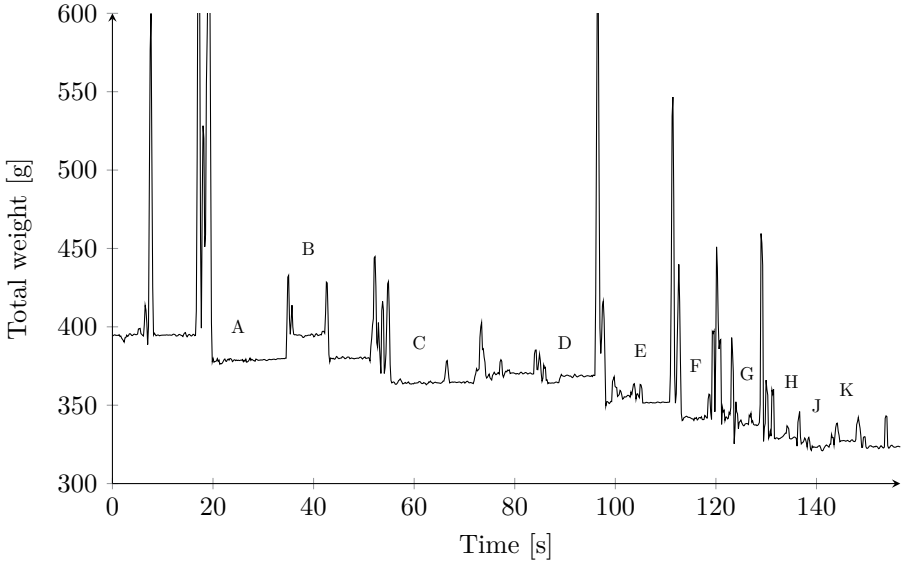
Notes:
[a] : C = Compartment. C1, C2 and C3 denote compartments 1 through 3.

7.3 Challenges During Unrestricted Eating

Restricted eating refers to the consumption of a meal, typically in a research or lab environment, where the subject is asked to change their regular eating behaviour. The subject may be asked to eat slower, limit their interaction with the food, refrain from cutting or re-arranging food, etc. Figure 7.2a shows a segment of plate data during restricted eating. In fact, the segment shown is the same data that was used in the preliminary study described in the previous section: it is the measurement taken in the lab environment with restrictions placed on the user. The total weight signal is shown. All labels correspond to a single bite. In this case, food was served pre-cut and the subject was instructed to limit their interaction with the plate (e.g.: no cutting). A bite is typically characterised by a large spike coming from scooping, pushing or stabbing into the plate to pick up the bite, followed by a decrease in weight and a period of stability where the subject brings the bite to the mouth. In this signal, every decrease in weight can be considered a bite.



(a) Restricted eating



(b) Unrestricted eating

Figure 7.2: Restricted eating versus unrestricted eating, captured with the smart plate system. In (a), every decrease in weight corresponds to a bite. In (b), labels C and D are partial bites while B and K are periods where the utensils are rested on the plate. The remaining labels are single bites.

Figure 7.2b shows an example of unrestricted eating. The signal appears a lot more *noisy* due to an increased amount of interaction with the plate and identifying bites is no longer straightforward. Bites are still characterised by the same pattern as before, but due to an increased eating speed and interaction with the plate, a decrease in weight no longer corresponds to a bite by default. During restricted eating, foods are typically served that can be eaten in a single bite without returning food to the plate. During unrestricted eating, however, the data may contain partial bites, where the subject returns a partially eaten piece of food back to the plate (such as pizza or a sandwich). In this example, labels C and D are partial bites, labels B and K are regions where utensils are temporarily placed on the plate and the remaining labels are single bites.

7.3.1 Bite Types

While social eating habits and customs can vary greatly between people, the underlying mechanics of consuming a meal remain the same. There are a few characteristic bite types that typically return during unrestricted eating, regardless of who is using the plate [29].

Figure 7.3 shows the *single bite*. This is the most common bite when eating with utensils. A single piece of food is picked up and consumed. Prior to t_1, the system may be unstable from a previous bite. Between t_1 and t_2, the weight remains stable as the user chews the previous bite. Then, the bite is picked up between t_2 and t_3 and the weight on the plate decreases. After the bite, the weight remains stable again as the user chews the bite. The weight of the bite is found by taking the difference in weight $w_1 - w_2$. The stable regions surrounding the bite are, however, not guaranteed. Immediately after picking up and consuming the bite, the user may already be manipulating the next bite, reducing the stable periods $t_1 - t_2$ and $t_3 - t_4$ or eliminating them altogether. This is one of the major challenges of bite detection during unrestricted eating as a lack of stable regions makes it difficult to determine the bite's weight.

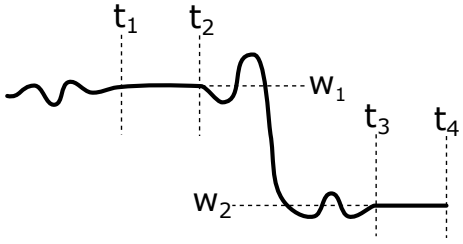


Figure 7.3: Hand drawn example of a single bite.

Figure 7.4 shows the *partial bite*. A partial bite occurs when the user takes a piece of food from the plate, takes a bite, and returns an uneaten part of food back to the plate. This type of bite is commonly observed when food is consumed with the hands, such as eating a slice of pizza or a sandwich. However, it does also occur when eating with utensils, for example when the user takes a bite out of a piece of meat and returns the uneaten part back to the plate.

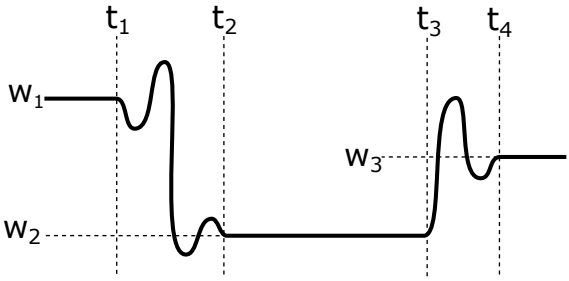
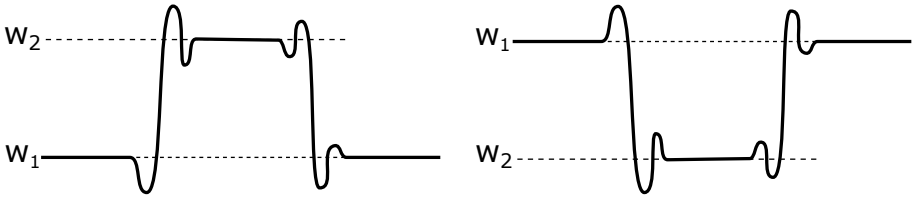


Figure 7.4: Hand drawn example of a partial bite.

The beginning of a partial bite is similar to a single bite. Between t_1 and t_2, a piece of food is removed from the plate. As the user takes a bite from the piece of food, the weight remains stable. Then, between t_3 and t_4 the unfinished piece of food is returned to the plate. The end weight w_3 is higher than the initial decrease to w_2, but lower than the start weight w_1. If this bite was treated as a single bite, its weight would have been incorrectly measured as $w_1 - w_2$, while the correct weight for this partial bite is $w_1 - w_3$. Care should thus be taken to differentiate between single and partial bites when measuring bites during unrestricted eating. Furthermore, multiple partial bites may happen in consecutive order. This is also observed in Figure 7.2b. Labels C and D show two partial bites in succession.

Another common type of occurrence during unrestricted eating is the activity of resting utensils on the plate for a while. This is illustrated in Figure 7.5a. The weight increases by the weight of the utensils and remains there as the user lets the utensils rest on the plate. As the utensils are removed, the weight signal returns to the same level prior to the introduction of the utensils. Label B in Figure 7.2b shows an example of this situation on real data measured with the smart plate during unrestricted eating.

The exact opposite of this situation is also possible. The user may take a food bite from the plate, but decide against eating it and return it to the plate in full. This is illustrated in Figure 7.5b. As the food bite is removed and placed back on the plate, the weight signal simply returns to the same level.



(a) Utensils resting (b) Zero bite

Figure 7.5: Hand drawn examples of utensils resting and zero bite.

7.4 Methods

As shown in the previous section, the challenges during unrestricted eating call for a more advanced kind of bite detection algorithm than the one proposed during the preliminary study. In the remaining sections of this chapter, a novel approach to bite detection is introduced. Machine learning algorithms are great for generalising concepts from large amounts of data, given a good feature set. Instead of manually defining rules and measures to detect a bite, we therefore propose the training of a machine learning algorithm with labelled smart plate data to automatically learn the patterns in the data that correspond to a bite.

7.4.1 Data

Data was collected in a nursing home in Belgium and a hospital in China. The PLATE-REMY dataset was collected from 14 older adults (avg. age 82 ± 10) in the nursing home. The PLATE-CHINA was collected from 18 older adults (avg. age 63 ± 10) in the hospital. In both cases, participants were healthy and could independently consume the meal. In all cases, food was consumed with utensils. In PLATE-REMY this was knife and fork, while for PLATE-CHINA food was consumed with chopsticks. The true bites were labelled in the data by a human observer. A total number of 317 and 836 true bite events were labelled in the PLATE-REMY and PLATE-CHINA datasets, respectively. Each true labelled bite is assigned a compartment number and the bite type (single or partial bite). Furthermore, the true amount of consumed food was recorded for each compartment on the smart plate. This was done by measuring the weights of food in each compartment with a kitchen scale before and after each meal. A more detailed overview of the data and the data collection process is given in Chapter 4.

7.4.2 Stability Detection

The bite detection algorithm is centred around the idea that every bite is surrounded by at least two stable regions: one prior to and one following the bite. The weight of the bite can then be measured by taking the difference in weight between both stable regions. This concept is also illustrated in the real plate data shown in Figure 7.2b. Labels A, E and F indicate single bites where a stable region can be observed before and after the bite. For label E, there is some noise right after the bite, but it is nevertheless followed by a stable region. For label F, the stable region after the bite is short and slightly noisy, but this can be filtered. The same concept applies to partial bites, as shown by labels C and D. Before a piece of food is removed from and put back on the plate, a stable region is observed. There are, however, exceptions to this. Labels G, H and J show single bites where the stable regions are not so visible any more. When a person is eating very fast or continuously manipulating the plate, bites may be inevitably lost. The bites G, H and J, however, may still be recoverable after filtering and processing of the raw data.

To determine a region of stability, a moving window with a length of 1 s is first passed over the total weight signal. The total weight signal is the sum of the three load cells after calibration. If the standard deviation of the data in the window is smaller than a threshold value, the window is considered to be stable. The threshold value was experimentally determined at 3 gram. In a second step, consecutive windows of stability are merged together, forming a single region of stability. The values within each window of stability are then replaced with the mean value over the entire region. Figure 7.6 shows the result of this process on a segment of data obtained during unrestricted eating. The total weight signal is shown. The thick horizontal black shaded areas are the detected regions of stability. The arrows indicate which pair of stable regions corresponds to a bite. The true labels labelled by an observer are shown with thick vertical lines. When a true label is found between two consecutive stable regions, that pair of stable regions is a bite. Note that annotation labels also contain the compartment number and the bite type (single/partial), but these are not shown in Figure 7.6.

7.4.3 Bite Classifier

Bites are detected by looking at the relative position of stable regions. As shown in Section 7.3.1, different bites or events have a different arrangement of consecutive stable regions. To detect a partial bite, for example, at least 3 consecutive stable regions are required to recognise the bite, as illustrated in Figure 7.4.

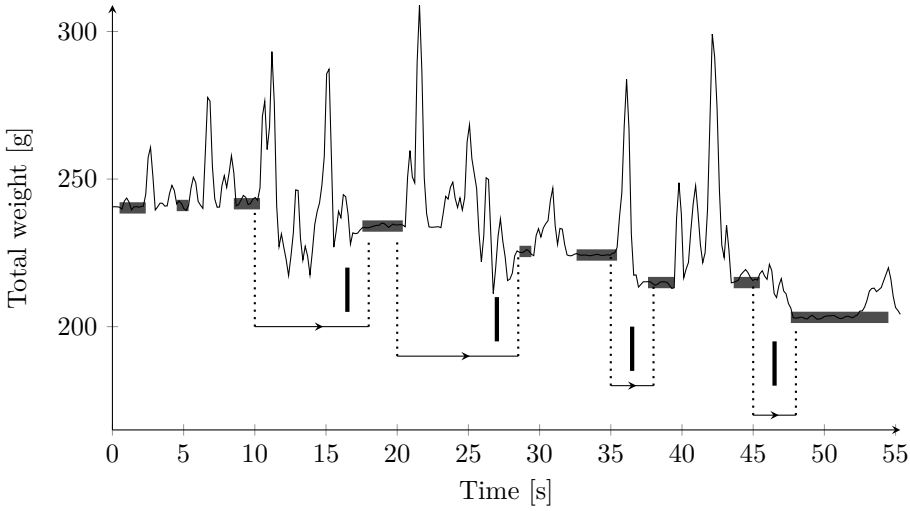


Figure 7.6: Plot showing the stable regions extracted from the raw total weight signal, denoted by black shaded areas. The thick vertical lines are the true annotations labelled by an observer. The dotted lines and arrows indicate which successions of stable regions the bite detection algorithm *should* detect as a bite.

Instances for the classifier are thus defined as groups of stable regions. The feature vector assigned to each instance looks at the changes in weight and time between adjacent stable regions. Given a total number of stable regions M, the weight and time information of each stable region is stored in a vector $\boldsymbol{S_i} = [W_i, T1_i, T2_i]$, with W_i the mean weight value of the region and $T1_i$ and $T2_i$ the respective start and stop times of the region, for $i = 1, ..., M$.

A feature vector is then constructed that combines information from 3 consecutive stable regions, defined by:

$$\boldsymbol{X_i} = \Big[(W_i - W_{i+1}), (W_i - W_{i+2}), (T2_i - T1_{i+1}), (T2_i - T1_{i+2})\Big]$$

for $i = 1, ..., M$. It combines the differences in weight between the first and second and between the first and thirds consecutive stable region. This is done for all stable regions in the signal. The complete feature matrix is thus written as $\boldsymbol{X} = [\boldsymbol{X_1}, \boldsymbol{X_2}, ..., \boldsymbol{X_M}]^{\mathrm{T}}$.

Constructing the class column vector $\boldsymbol{Y}$ is done based on the true bite labels assigned during the video-annotation. True bite labels were always placed at the time immediately after a bite was removed from the plate. In the case of a

partial bite, it was placed in the beginning where the bite was first removed from the plate. Class labels are thus assigned based on the presence of a class label between the first and second stable region.

If a single true bite label is located between $T2_i$ and $T2_{i+1}$, the instance $\boldsymbol{X_i}$ is labelled as a bite. The stop time, instead of the start time, of the second stable region is used to allow for some inaccuracy in the labelling. If no true label is found in this region, the instance is labelled *nobite*.

A special situation occurs if multiple true labels are found between $T2_i$ and $T2_{i+1}$. In this case, the instance is not unequivocally assignable to either class. This case can be called a *multi-bite*. It happens when no stable region was found in-between two or more true bites, which can happen, for example, when the user is eating fast. The primary goal of the classifier is to detect individual bites so that they can be measured, localised and assigned to a compartment in order to detect the total consumed weights per compartment. In the event of a multi-bite, the individual true bites are no longer detectable and cannot be localised on the plate. It is thus reasonable to assign a multi-bite to the negative class of the classifier and assign it a *nobite* label. Doing so, however, has the potential to introduce noise in the classifier. A multi-bite corresponds to a decrease in weight, similar to a single bite. This can bias the classifier to miss single-bites with a similar decrease in weight.

To evaluate the impact of multi-bites, two test cases are evaluated. In *test case 1* (TC1), multi-bite instances are assigned the *nobite* label during training, in an attempt to train a classifier that only detects single bite events. In *test case 2* (TC2), multi-bites instances are assigned the *bite* label. This alternative classifier detects both single- and multi-bite events, effectively creating a classifier that detects epochs that correspond to eating behaviour in general.

For both test cases, a supervised classifier of the Random Forest (RF) type is trained with feature matrix $\boldsymbol{X}$ and class labels $\boldsymbol{Y}$ to distinguish between bite and no-bite instances. The RF is chosen for its ability to generalise eating behaviour of multiple subjects in a single model. Eating behaviour can vary between test subjects, so a single decision tree may not be generalised enough to describe the eating activity of different test subjects. The RF trains multiple trees, increasing the performance of the model, as described in Section 2.3.5. The classifier makes no distinction between single or partial bites. A distinction is only made when determining the bite weight. The weight of the detected bite BW_i is calculated by looking at the differences in weight of W_i, W_{i+1} and W_{i+2}, with following formula:

$$\mathrm{BW}_i = \begin{cases} W_i - W_{i+1} & \text{if } (W_i > W_{i+1} > W_{i+2}) \ \rightarrow \text{single bite} \\ W_i - W_{i+2} & \text{if } (W_i > W_{i+1} < W_{i+2}) \ \rightarrow \text{partial bite} \end{cases} \tag{7.1}$$

7.4.4 Detection Pipeline

Figure 7.7 shows the complete algorithm pipeline. The stability detection block takes the total weight signal w_{tot} and looks for stable regions $\boldsymbol{S}$. After finding the total number of stable regions M, for $i = 1, ..., M$, the stable region $\boldsymbol{S}_i$ and its two subsequent neighbours, $\boldsymbol{S}_{i+1}$ and $\boldsymbol{S}_{i+2}$, are passed to the feature detection block. The feature extraction block outputs the feature vector $\boldsymbol{X}_i$, as described in the previous section. The feature vector is then classified by the bite detection model. The prediction $\hat{Y}_i$ is passed to the compartment detection. The compartment detection attaches a compartment number to each positive detection. Negative detections are passed through with no change, because the bite detector determined there is no bite to localise. The compartment detection takes the total weight of the bite as well as the weights measured by each individual load cell and localises the bite using the bite localisation formula presented in Section 6.2. Total weight of the bite is calculated using Formula 7.1. The coordinates are then compared to the layout of the polymer plate and a compartment number from 1 to 3 is assigned to each detection. At the output of the algorithm, each instance will have been assigned either a compartment number or a *nobite* label.

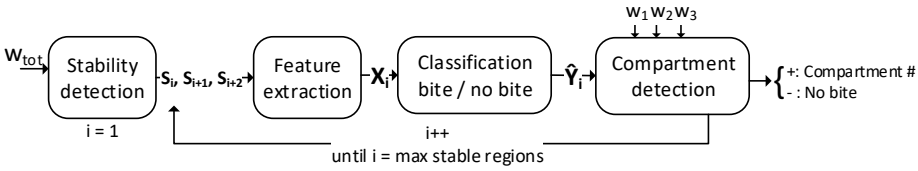


Figure 7.7: The overall detection pipeline.

7.4.5 Performance Metrics

The bite detection classifier is evaluated using a leave-one-person-out (LOPO) cross-validation scheme. Data from one person is used to test the algorithm, and data from all other persons is used to train the classification model. This is done for each person in the dataset. Note that only the bite/nobite classifier is retrained in each iteration, the compartment detection is a deterministic model that remains unchanged.

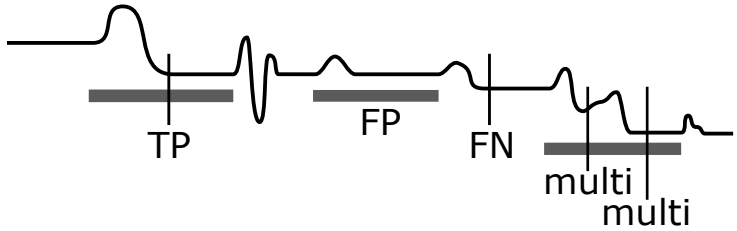


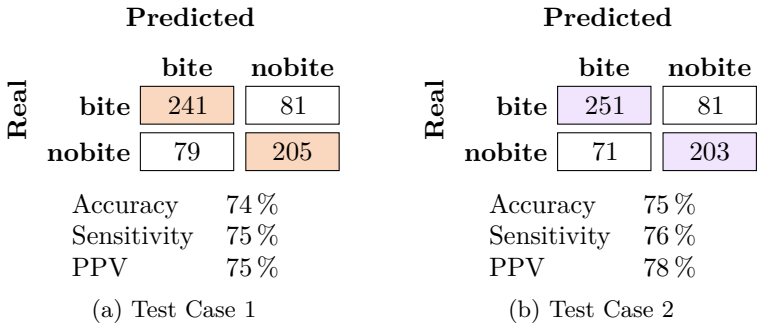
Figure 7.8: Hand drawn example of the classification results. Thick horizontal bars indicate a detection, thin vertical lines are the true labels.

To evaluate the algorithm's performance for each test iteration, the output labels are compared to the ground truth labels obtained during annotation. For the evaluation, no distinction is made between single and partial bites. This information is only used when weighing the individual bites. Evaluation is done in a similar fashion as the labelling of the training data. The predicted bites contain the beginning and end times of a bite, their weight and their predicted location. These are then compared to the annotated ground truth labels. If one true label occurs inside the start and stop time of a predicted bite, this is considered a True Positive (TP). If no true label is found inside the detection range, the bite is considered a False Positive (FP). If multiple true labels occur within one predicted bite, this is a multi-bite and its evaluation depends on the test case. In test case 1, where the classifier was trained to detect only single-bite events and multi-bites were assigned the negative class during training, this is considered a FP. In test case 2, where multi-bite instances were assigned the positive class during training, the prediction is considered a TP. If a true label is found in a region marked as negative (nobite), this is a False Negative (FN). If no true label is found in this region, it is a True Negative (TN).

Figure 7.8 shows the evaluation methodology on a hand drawn example signal. The bite classifier outputs regions during which it has detected a bite (the end times of the first and second stable segment). The thick horizontal bars show regions where the classifier has detected a bite. True labels are indicated with thin vertical lines. An evaluation is made based on the position of the true labels relative to the detected regions.

Furthermore, the positive bite detections are used to estimate the weight of consumed food from each compartment. This is done for each person individually. The consumed food per compartment is obtained by taking the sum of weights of bites detected in each compartment. This value is then compared to the true weight of consumed food and the error per compartment is calculated. Recall that the true weight of consumed food is known because the food in each compartment was measured with a kitchen scale before and after each meal.

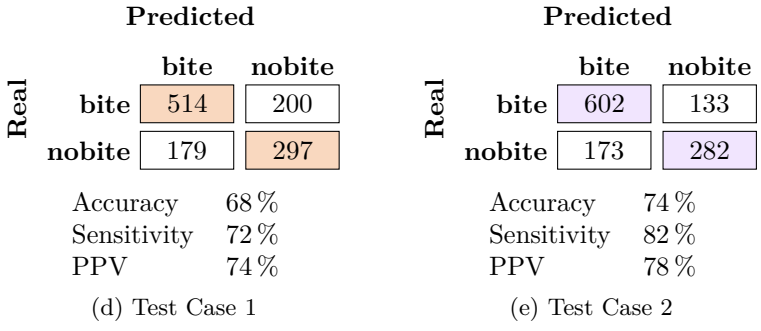
Table 7.2: Confusion Matrices: bite detection.



(a) Test Case 1

(b) Test Case 2

(c) PLATE-REMY dataset



(d) Test Case 1

(e) Test Case 2

(f) PLATE-CHINA dataset

7.5 Results

Table 7.2 shows the confusion matrices of the bite detection for the PLATE-REMY and PLATE-CHINA datasets. The results of both test cases are shown. The detection accuracy, sensitivity and Positive Predictive Value (PPV) are indicated below the confusion matrix for each test case. For the PLATE-REMY dataset, there is little difference between both test cases. Test case 2 performs marginally better, with a PPV and sensitivity of 78 % and 76 % respectively. For the PLATE-CHINA dataset, there is a significant difference between test cases. This is because the PLATE-CHINA dataset contains approximately 7 times more multi-bite events. For the PLATE-CHINA dataset, test case 2 performs the best with a PPV and sensitivity of 78 % and 82 % respectively. In test case 1, the classifier detected 38 multi-bites that are included in the

179 FP events, while 43 out of 602 TP are multi-bites in test case 2. This shows that, with the exception of 5 instances, the classifier is not able to classify multi-bites separately from single-bites, as they are detected as a bite in both test cases. The marginal decrease in detected multi-bites in test case 1 does not outweigh the increase in performance in test case 2, particularly in the number of FN events. This supports the hypothesis that multi-bites events are not distinguishable from single-bite events and assigning them the *nobite* label during training has a negative impact on classifier performance.

Table 7.3 shows the confusion matrix when taking into account the predicted compartment numbers for each bite. Shown are the results on the PLATE-CHINA dataset. The positive detections are now split up according to their predicted and real compartment. The FP and FN are still shown under the *nobite* label, but are also split up based on their predicted and real compartment respectively. This is the result of FPs that were wrongly classified as being a bite, and thus were assigned a wrong compartment because there was no bite in the first place, or FNs that should have been assigned a compartment, but were missed because the compartment detection only attempts to localise the positive detections of the classifier. The numbers between brackets are the amount of additional multi-bite detections that contain true bites from the same compartment. In this situation, the compartment of the true bites is assigned as real compartment of the multi-bite. There are 26 multi-bites, not included in the confusion matrix, where a real compartment could not be assigned because the true bites came from different compartments.

Figure 7.9 shows a box plot of the error per compartment over all measurements for the PLATE-REMY and PLATE-CHINA datasets. The error for each measurement is normalised to the total weight of each meal. A negative error means that the algorithm under-reported the actual food intake per compartment. A positive error is the result of over-reporting. The PLATE-REMY dataset shows slightly less spread and less outliers than PLATE-CHINA. This can be attributed to the fact that PLATE-REMY is a smaller dataset, so the dataset may be biased, while PLATE-CHINA may be more generalised and representative for unrestricted eating. For the PLATE-CHINA dataset, the results are more biased towards a negative error, which can be explained by the fact that the bite detection reported more FN than FP. Similar to the results of the bite detection algorithm, test case 2 performs the best with lower maximum errors, lower median error and a lower interquartile range (IQR). The average error and std. dev. per compartment for test case 2 is $(12 \pm 12)\,\%$, $(5 \pm 5)\,\%$ and $(8 \pm 9)\,\%$ for compartment 1, 2 and 3 respectively. This is the error relative to the total portion size. Looking at the total weight of the meal, the algorithm is able to estimate the total portion size with an average error of $(13 \pm 12)\,\%$.

Table 7.3: Confusion Matrix: bite localisation for the PLATE-CHINA dataset.

Predicted compartment

Real compartment	**1**	**2**	**3**	**nobite**
1	60	19	16	33
2	9	96	14	59
3	17	42	241	108
nobite	26	63	90	297

(a) Case 1

Predicted compartment

Real compartment	**1**	**2**	**3**	**nobite**
1	62	20	18	22
2	10	105 (3)	15 (2)	41
3	22	43 (2)	264 (10)	70
nobite	30	58	85	282

(b) Case 2

Note 1: because the localisation algorithm can only localise the positive outputs of the classifier, it will also attempt to localise false positive bites and miss false negatives. These are shown under the *nobite* label. True negative bites are shown on the crossing of the *nobite* label.

Note 2: the numbers between brackets indicate the amount of additional multi-bite detections that contain true bites all from the same compartment. In this case, the compartment of the true bites is assigned as the *real* compartment of the multi-bite. There are 26 remaining multi-bites where the true bites inside the multi-bite event come from different compartments. The latter are not included in the confusion matrix because no *real* compartment can be unequivocally assigned.

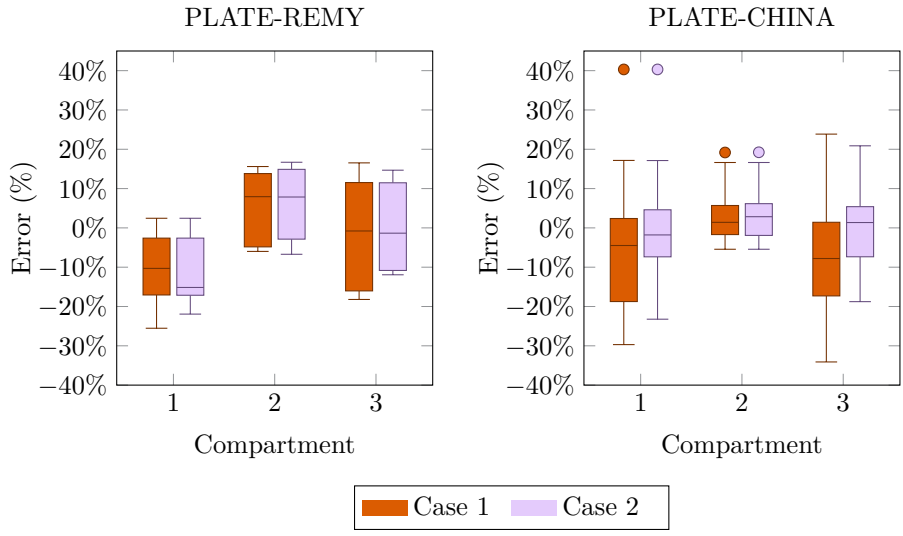


Figure 7.9: Box plot showing the distribution of the weight error per compartment, relative to the total weight of the meal. Shown for both the PLATE-REMY and PLATE-CHINA datasets.

7.6 Discussion

The biggest challenge during unrestricted eating is dealing with a variety of different situations that can arise when a subject interacts with the plate and food. Without additional context information, it can be difficult, or sometimes impossible, to extract individual bites from the weight signal. One factor contributing to misclassifications are bites taken in quick succession without allowing the plate to stabilise and the re-arranging of food. Almost all of the measurements in our dataset contained segments of data with multiple bites, but with no distinguishable stable regions in-between, making the identification of single bites impossible. In this situation, however, the cumulative size of the bites could often still be determined. Another source of errors is the misclassification of partial bites as single bites. If the weight signal is not stable while a partial bite is being taken, for example because the subject is using their off-hand to touch or lean on the plate, the replacement of the partial bite on the plate may be missed and the initial decrease in weight may be seen as a single bite. Some subjects had the habit of touching or leaning on the plate with their free hand while taking a bite, again preventing the plate from stabilising.

An interesting difference could be observed in both datasets that is related to

cultural differences. In the PLATE-CHINA data, a frequent occurrence is the re-arranging of food in order to pick up the bite. This is directly related to the use of chopsticks. Subjects would often poke at or shove pieces of food to align them in an orientation that made them easier to pick up with chopsticks. In some cases, a single bite of food would be picked up multiple times from the plate and put back in order to get a better grip on it, before being consumed. This in turn is a source of some false positives. This activity was less pronounced in the PLATE-REMY data, where subjects typically scooped up or stabbed into the food with a more fluent movement.

Naturally, misclassifications will also result in an incorrect weight measurement. While the median weight error per compartment is situated around 0 for the PLATE-CHINA dataset and the boundaries of the interquartile range (IQR) are less than 10 % of the total portion size for test case 2, as shown in Figure 7.9, outliers of up to 40 % are observed. The small IQR gives an indication that the algorithm performs well for a majority of the test data, while the large minima and maxima and few outliers indicate that there are still some measurements where the algorithm performs subpar. As discusses earlier, these are situations where the algorithm was unable to sufficiently detect individual bites in each compartment. For example, all outliers observed in the box plot are from a single measurement where the test subject was touching or leaning on the plate for long periods of time. Similarly, the large maximum values for compartment 1 or 3 are also the result of another test subject leaning on the plate for over half the duration of the meal.

Even if the detected weight per compartment is incorrect, however, the total meal size may still be correctly measured. A bite may be classified in the wrong compartment, but its weight can still be measured correctly and added to the estimated total portion size. In some cases, bites that were taken in quick succession were not individually detected, but their cumulative weight decrease was detected as one large bite, again contributing to a correct estimation of the total portion size. On average for both datasets, the algorithm was able to estimate the total portion size with an error of 13 %. If we disregard the measurements where test subjects leaned on the plate for large periods of time, an activity that makes it almost impossible for the bite detection algorithm to perform well, the average error of the total portion size drops to 9 %. One study found that nursing staff, on average, overestimate the amount of food consumed in nursing home residents by 22 % [103]. An automated plate system could further help reduce work load and provide more accurate intake estimation for nursing home residents.

In some cases, the algorithm performed exceptionally well. In 9 measurements of the PLATE-CHINA dataset, the algorithm measured the weight of the food consumed in each compartment with an error of less than 15 gram. In one case,

the meal consisted of dumplings that were served only in compartment 3. In this measurement, the algorithm approached perfect accuracy, with no bites detected in compartment 1 and 2, 96 % of bites detected in compartment 3 and with the total weight correctly estimated within 9 gram.

Related work by Mattfeld et al. reported a correct identification of 39 % of true bites and drink events (out of 24101) using a table embedded scale, tested on a dataset of 271 subjects [29]. While our approach is able to detect 72 % of bites out of 836 for PLATE-CHINA and 79 % of bites out of 317 for PLATE-REMY, it cannot measure drink events and our dataset is much smaller. Our proposed plate system does offer the advantage of being mobile and has the ability to estimate the location of each bite.

7.7 Strengths and Limitations

The results show that the presented algorithm together with the smart plate is capable of measuring bites during unrestricted eating with a high accuracy. Some situations remain, however, where the algorithm is not able to measure a bite. Because the algorithm relies on the presence of stable regions in the weight signal, activities that disturb these stable regions will have a negative impact on performance. The activity of leaning or resting a hand on the plate was observed several times in the captured data, which causes the weight signal to 'drift' and prevents it from stabilising. If this happens when bites are taken from the plate, the bites are typically not measurable. This is one of the limitations of the current system. In the future, this could be solved by adding an additional enclosure around the plate that is not in contact with the weight sensors, allowing users to lean or rest a hand on the plate without influencing the weight signal. This would of course increase complexity of the system and the lack thereof is currently one of the major strengths of the smart plate. The hardware is relatively cheap and the system can be made to operate in the background, making it easy to use by older adults. Furthermore, this is the first time that a method is presented that has the potential of measuring the consumption of different food types from a single plate without embedding sensors into the plate itself.

Our work shows that the measurement of consumed food per compartment is possible, but the accuracy of the weight detection is directly dependant on the performance of the bite detection algorithm. Only when bites can accurately be detected in the weight signal, can the consumed food per compartment be measured. In half of the measurements in the dataset, food was measured with an error of less than 15 g, which, on a total meal size of 200 g to 500 g, is low.

It remains to be seen, however, how the smart plate performs when used for a systematic monitoring of food intake. Currently, its accuracy has not yet been compared to one of the ground truth methods for the monitoring of food intake. This will require additional validation studies. In order for the smart plate to be successful, the measurement error should at least be in the same order as that of the ground truth method.

7.8 Conclusion

A novel approach to bite detection has been presented in this chapter. The overall smart plate system with bite detector and bite localisation shows promising results. The system is able to estimate where bites are taken on a plate. With prior knowledge of food types served in each compartment, the total amount of ingested calories can be estimated. Until now, detecting individual food types with a weighing system in this manner has not been demonstrated in the literature. The concept of localising bites or events on a weighing surface may provide future perspectives for food monitoring applications.

Chapter 8

Quantifying Eating Behaviour With a Smart Plate

This chapter presents an exploratory study that proposes the extraction of behavioural parameters from the smart plate data. These parameters can be used in the quantification of eating behaviour and may be used for long-term monitoring of eating behaviour.

8.1 Introduction

The work presented in the thesis so far has focussed on the measurement of eating activity from a perspective of quantifying the raw food intake. Using wearable sensors we can measure when or how long a person is eating, while the smart plate measures the quantity of consumed food. These parameters can provide an objective overview of the food intake that can be stored in a food diary, to be consulted by a dietician or physician.

These parameters alone, however, may not give a complete overview of the food intake. The behaviour of subjects during the eating activity is an equally important parameter to consider [49, 50, 104, 105]. Eating behaviour can differ between persons, e.g. speed, usage of utensils, interaction with the food, timings, etc. This information can provide more contextual information about individual meals or food intake over time. These measures may give an indication about the physical or emotional state of the subject. Changes in behaviour may indicate a problem in the health status of the subject. For example, a decreasing

bite size over time may indicate a problem with picking up food, indicating that a change in utensils or food type is required. Research has shown that eating behaviour is largely habitual in nature [104]. A subject drifting from these habitual clues can be an indicator for psychological or physical change in the subject, potentially requiring an intervention.

Eating behaviour has been widely studied in the literature. The gold standard for quantifying eating behaviour is the use of questionnaires [105–107]. These questionnaires potentially suffer from the same shortcomings as questionnaires for the quantification of food intake: they are time intensive to use for longer periods of time and may contain mistakes due to human error. With the introduction of sensor based automatic food intake monitoring techniques, these kind of questionnaires can be supplemented with data recorded by the sensors.

In the remainder of this chapter, a number of parameters that can be extracted from the smart plate data are proposed and illustrated with exploratory results.

8.2 Bite Parameters

The previous chapter saw the introduction of the smart plate system and corresponding bite detection algorithm. A measurement of the food intake down to the individual bite level has been proposed in the literature as a method for gaining more insight in eating behaviour and changes in environmental clues or health status [29, 46–50]. Until now, individual bites have remained largely undetected in eating data during unrestricted eating [29].

The presented smart plate system and bite detection algorithm can provide new insights into eating behaviour. Parameters can be extracted from the individually detected bites that can give more information compared to the total amount of food intake. Individual bites can be a measure of the microstructure of eating behaviour [46–48]. These are characteristics of food intake that can be highly personalised. While the overall activity of eating is similar for all humans, there can be substantial differences in personal preference, food types or cultural and social habits. We therefore propose the extraction of parameters from the bites detected with the smart plate. A selection of parameters that can be extracted is given below.

- Bite size
- Bite length
- Time between bites
- Total bites
- Bites per minute
- Bites per compartment
- Bite pickup time (time to place food on utensil and pick it up)

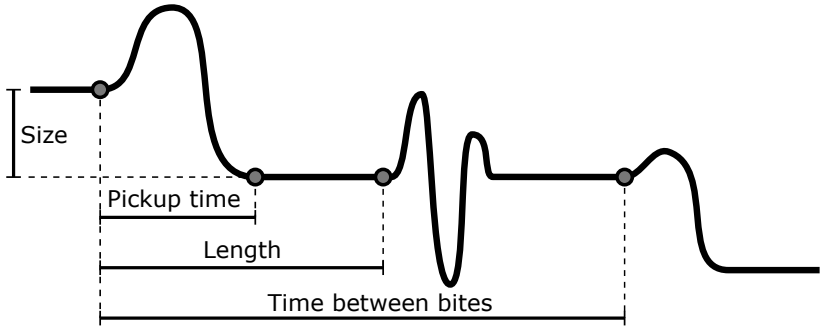


Figure 8.1: Bite parameters illustrated on a hand drawn weight signal.

Figure 8.1 illustrates how bite parameters can be extracted on a hand drawn weight signal with two bites. These parameters can be directly derived from the output of the bite detection algorithm presented in Chapter 7. Bite pickup time is the time required for a piece of food to be placed on the utensil, while bite length is the time from when the user begins taking a bite until the end of the subsequent stable region. Bite length can be a measure of eating speed: a slow eater is expected to have a longer bite length.

As an example, Figure 8.2 shows the probability density function of bite length for two persons in the PLATE-CHINA dataset. Two subjects with a high performance for the bite detection algorithm were selected, with an F_1 score of above 85 %, as a low performance will also result in erroneous parameters. Both subjects have a different distribution, indicating that their behaviour during the meal, in this case eating speed, was different.

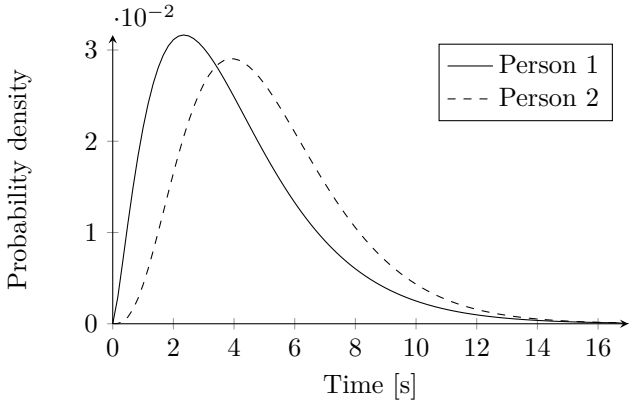


Figure 8.2: Probability density function of the bite length for two subjects in PLATE-CHINA. It shows the distribution for a single meal per person.

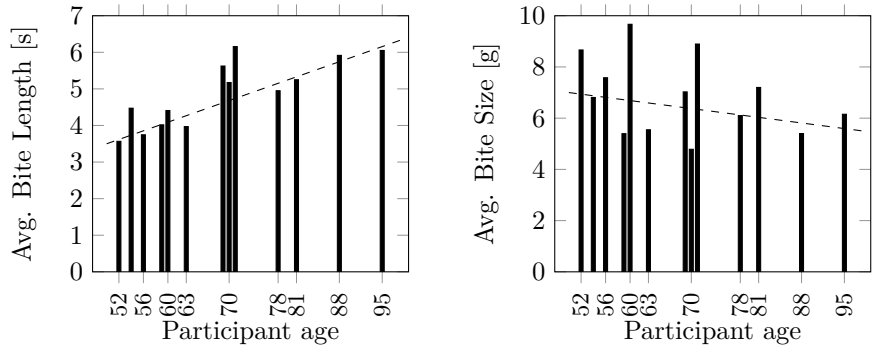
8.3 Age Specific Behaviour

In this section, we present exploratory results for age-specific food intake parameters extracted from the PLATE-CHINA dataset. While the focus of this thesis is food intake in older adults, several younger adults were included in the PLATE-CHINA dataset to get a more generalised dataset of eating activity. Out of the 18 subjects in the dataset, 7 are under the age of 60, but above the age of 52. As a result of this, the age differences in this dataset may provide some interesting insights into age-related eating behaviour as subjects increase in age.

Figure 8.3a shows a bar graph plotting the average bite length during the meal in function of participant age. In the event that multiple participants have the same age, the average value is taken. The bite length is defined as the time between the beginning of a bite and the end of the subsequent stable period. It is a parameter that can be used as a measure of eating speed. Fast eating will result in a low bite length. An increasing trend can be observed in the data, indicating that eating speed decreases with age. Furthermore, two vertical levels can be observed in the data. Subjects below the age of 70 all have a similar bite length, while subjects over 70 share an increased level of bite length. This is in line with what is already known about food intake in older adults [6, 95, 108]. Bite length may increase as a result of a decrease in muscle strength, trouble chewing the food or a problem with swallowing. It is a confirmation of the smart plate's ability to accurately extract individual bites during eating. This allows bite length to be monitored over time and used as an indicator of abnormal behaviour, as bite length is expected to stay relatively constant over time within age groups.

Figure 8.3b shows a similar bar graph with average bite size during the meal in function of participant age. For bite size, a slight downwards trend can be observed with age, but it is less pronounced and it is more difficult to make a conclusive link between both variables. The downwards trend may indicate that a decrease in muscle strength with age results in the inability to pick up larger amounts of food, but the difference is only a few grams between the youngest and older adults. The difference may be largely related to different food types. Since the PLATE-CHINA measurement only consists of one measurement per person, the food type can introduce a strong bias in the observed bite size. For a more conclusive indication on the effect of age on bite size, repeated measurements with different food types are required.

Figure 8.4 shows a bar graph showing the total bite count per meal in function of participant age. No conclusive observation can be made from these results. Bite count appears to vary widely over all age groups. This may again be related



(a) Avg. bite length per age

(b) Avg. bite size per age

Figure 8.3: Average bite parameters sorted by age in PLATE-CHINA.

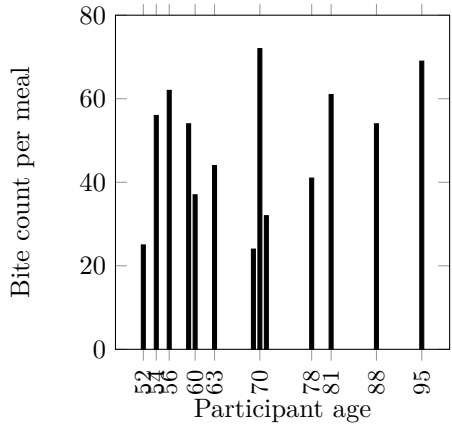


Figure 8.4: Total bite count per meal sorted by age in PLATE-CHINA.

to the type of food that is being served: some foods inherently require less bites to consume. Bite count is also related to the total amount of consumed food. Higher amounts of food intake should typically correspond to a higher bite count. Bite count could still be an interesting parameter to consider. When a person exhibits a small amount of consumed food in combination with a high bite count, this may indicate a problem with the physical activity of eating, such as picking up food or chewing.

While the findings presented in this section are preliminary and are most likely biased by food type or circumstantial events, they are in line with previously

reported findings in the literature that food intake changes with increasing age [1–4, 6–8, 95, 108]. Parameters such as bite length, size or count may be used to categorise older adults into groups of expected ranges. Given a certain age group and level of independence of the subject, bite length, for example, is expected to be in a certain range. Should a person deviate from this range, it can be an indicator of a behavioural or physical change. These changes need not be pronounced, this change may happen slowly over time, unnoticed by the subject or the caregivers.

8.4 Person Specific Behaviour

It is difficult to extract person specific behavioural parameters from the datasets that have been used in the previous parts of the thesis. To show person specific behaviour with any kind of certainty, repeat measurements of the same person are required to decorrelate the food intake from the served food type or circumstantial events during the meal. Furthermore, it is not straightforward to link one particular behavioural trait to the bite parameters without a reliable ground truth. For this purpose, an additional exploratory dataset was recorded in a hospital. Bite parameters are extracted from the dataset and linked to a known behavioural trait.

8.4.1 Data & Methods

Data was collected at the Department of Rehabilitation Medicine at Huashan hospital Jing'An Branch, affiliated to Fudan University. Six patients were selected that recently suffered a stroke. All patients had an upper limb defect as a result of the stroke, impairing their ability to fluently consume food. Patients were admitted to the rehabilitation department to regain control over their fine motor skills. Repeat measurements were taken from this group. Over the course of 15 days, patients consumed a meal with the smart plate system every 3 days, resulting in a total of 5 measurements per patient.

Table 8.1 gives an overview of patient characteristics. The impairment column indicates a level of impairment for the limb that was used to consume the food. The level of impairment was determined with the Fugl-Meyer assessment (FMA) performed by a doctor [109, 110]. The FMA is used to evaluate motor function after stroke. The FMA consists of a standardised set of exercises that are scored based on the physical ability of the patient. Full points (2) are given if the patient can complete the exercise, half points (1) if the exercise can only be completed partially and zero points if the patient is unable to perform the

Table 8.1: Participant characteristics of the PLATE-REHAB dataset.

id	age	impairment[a]	utensil[b]
1	54	none	chopsticks
2	33	severe	spoon
3	84	small	chopsticks
4	66	severe	spoon
5	77	small	chopsticks
6	46	severe	chopsticks

Notes:
[a] : indicates impairment of the upper limb used for bringing food to the mouth.
[b] : the indicated utensil was used for all 5 measurements of that patient

exercise. The complete FMA evaluates motor function of the entire body. For this study, only exercises related to the upper extremities are taken into account. The maximum score obtainable for the upper limb exercises is 66. Level of impairment is ranked as follows: < 32 severe, $32 - 50$ high, $51 - 61$ medium, $62 - 66$ small.

The difficulty to consume food due to the upper limb impairment can be seen as a behavioural trait, albeit unwanted. Furthermore, the level of difficulty will change according to the level of impairment. Because the level of impairment is known, we can attempt to link this behavioural trait to the parameters extracted from the smart plate data. For patient number 1 in the dataset, the defected limb was too severely impaired to consume food, resulting in this patient consuming food with his other, unimpaired, limb. This person can act as a control measurement, since food consumption was more fluent in this patient compared to all other patients.

Because this dataset is unlabelled, the entire PLATE-CHINA dataset is used to train the bite detection classifier. This model is then used to detect bites in the rehabilitation dataset. The same methodology as presented in Chapter 7 is used. While the lack of labelling makes verification of the bites impossible, the bite-related parameters can still provide interesting insights.

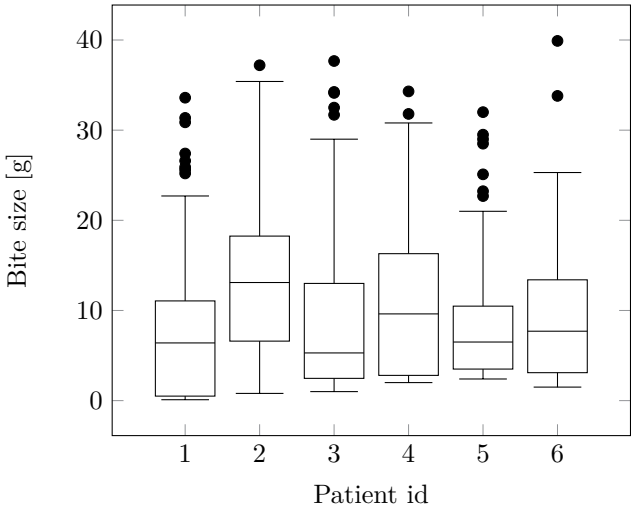


Figure 8.5: Box plot showing the distribution of bite size for each person in the PLATE-REHAB dataset. Patients 2 and 4 are the only patients in the dataset using spoons, as opposed to chopsticks. A higher median, interquartile range and upper whiskers can be observed for these patients.

8.4.2 Results & Discussion

Figure 8.5 shows a box and whisker plot containing the distribution of bite sizes per patient. Similarly to the age related behaviour discussed in the previous section, it is difficult to correlate any behavioural parameter to the bite size. There is, however, one observation that can be made. Patients 2 and 4 are the only patients in the dataset that consumed their meals with a spoon, this due to a difficulty with using chopsticks. The remaining patients all used chopsticks. As shown in the box plot, patients 2 and 4 have a noticeable higher median bite size, interquartile range and upper whiskers than the others. This indicates that bite size was consistently larger for these patients, which corresponds to the usage of a spoon. The spoons used were full sized table spoons. More food can be scooped in this kind of spoon than can be picked up with chopsticks. While this is, again, highly related to the food type, the results indicate that the type of utensil can be detected solely from the smart plate data. This, in turn, may provide further insights into the behavioural patterns of smart plate users.

Figure 8.6 shows the distribution of bite length over all participants. Bite length is defined as the time from when the bite starts to the end of the subsequent stable region in the smart plate data. It is a measure of eating

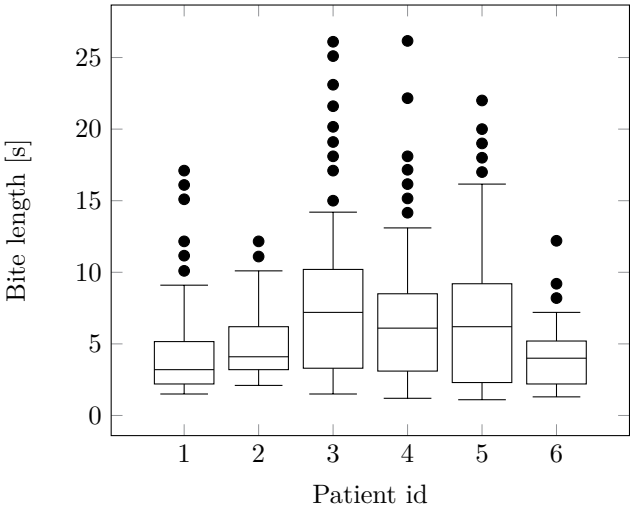


Figure 8.6: Box plot showing the distribution of bite length for each person in the PLATE-REHAB dataset. Patients 3 through 5 are the only patients over the age of 65. Their increased spread in the distribution may indicate that bite length is proportionate with age.

speed, as eating fast will result in the user immediately picking up bite after bite, shortening the time of the subsequent stable region. Eating speed is thus inversely proportionate to bite length. In the previous section, we proposed that bite length may increase with age. The only three patients in our dataset over the age of 65, patients 3-5, exhibit a consistently larger bite length than the other three patients. Furthermore, patient 4 had a *severe* level of impairment, which may also contribute to the larger bite length. Patients 2 and 6, however, also had a *severe* level of impairment and do not exhibit a noticeably large bite length. Bite length may not be a definitive descriptor of impairment, as an impaired patient might still attempt to eat fast.

Figure 8.7 shows the distribution of time required to pick up a bite, also referred to as bite pickup time. Bite pickup time is defined as the time the user first starts interacting with the plate in an attempt to place a bite of food on their utensils, to the time the bite is actually picked up from the plate. For impaired patients, bite pickup time is expected to be long, as the decrease in fine motor skills makes it difficult to pick up food from the plate. They may require multiple attempts to pick up the bite, increasing the bite pickup time. This is also one of the reasons patients 2 and 4 used spoons, which make it easier to pick up food. Patients 4 and 6 with *severe* impairment show an increased median pickup time, high spread and high upper quartile value. This indicates

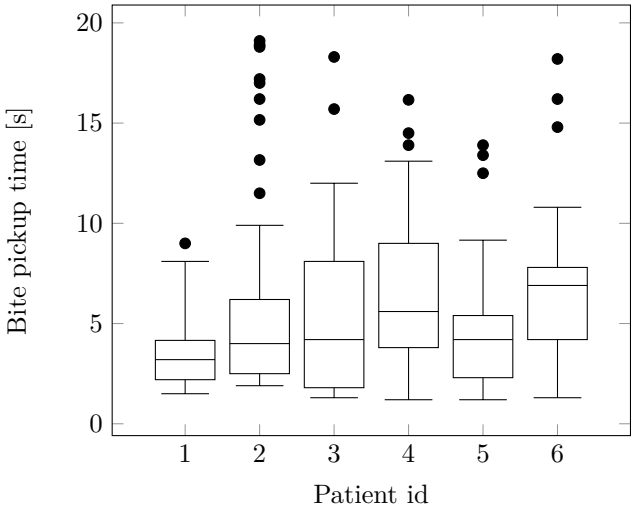


Figure 8.7: Box plot showing the distribution of the time required to pick up a bite in the PLATE-REHAB dataset. Patients 4 and 6 with severe impaired exhibit the highest median pickup times and lower quartile range. This may confirm that bite pickup time is an indicator of impairment.

that these patients had the most difficulty pickup up bites, which correlates to their impairment status. Patient 2 with *severe* impairment has a lower median pickup time and spread, but does show considerable outliers upwards of 10 s. Patient 3 with *small* impairment shows a high interquartile range and high upper whisker, but this is coincidentally also the oldest patient in the dataset, so high pickup time in this patient may be due to age rather than impairment. The non-impaired patient number 1, shows the smallest median and spread of all patients, further adding to the assumption that there may be a correlation between bite pickup time and impairment.

8.5 Conclusion

This chapter has introduced the concept of using bite parameters extracted from the smart plate data as descriptors of behavioural traits. A dataset of patients recovering after stroke with upper body impairment was used to show that there may be a correlation between these parameters and levels of impairment, which can be seen as a behavioural trait. The results presented in this chapter are a first step to further studies with the smart plate for the study of behavioural traits. In future work, the smart plate may also be clinically evaluated for patient monitoring and behavioural analysis in patients with different pathologies.

Chapter 9

Conclusion & Future Work

9.1 Summary of Results

The goal of this thesis was to investigate modalities for food intake monitoring in older adults. Traditional monitoring tools such as questionnaires are inherently challenging to use in an older population. Sensor based automatic food intake monitoring may provide an alternative. These systems can be made to work in the background and if implemented well, can provide better and more objective measurements of the daily food intake. Research and state of the art in this area is primarily aimed at young to middle aged adults, and research specifically aimed towards older adults is sparse to non-existing. Older adults, however, are at risk for malnutrition, one of the conditions that typically require an intake monitoring study. With an increasingly growing ageing population and limited resources, automatic intake monitoring may play an important role in the routine care of older adults. The tools presented in the thesis aim to assist in or partially replace traditional self-reporting and screening tools. By lowering the recording burden on the user, continuous food intake monitoring can be made easier and more comfortable to use. This can make recognition and prevention of malnutrition more accessible in situations where traditional methods fall short, such as an older adults living alone at home.

The thesis presented a study into the use of an accelerometer for the measurement of chewing activity. Detection of chewing activity has been studied in the literature as a descriptor of eating activity. Existing methods, however, typically rely on uncomfortable or obtrusive sensors. While the detection of movement of mastication muscles has been used as an indicator of chewing, accelerometers

have not been used yet for this task. Accelerometers are a versatile sensor that may be used for the detection of other eating related activities, such as drinking, preparing food, cooking, etc, in addition to being used for the detection of chewing. The method proposed in this thesis utilises an accelerometer mounted to the eyeglasses of the user. Eyeglasses are used in this context as a way of mounting the accelerometer in such a way that it makes indirect contact with the mastication muscles. The ears of the glasses are able to transmit the movements and vibrations during chewing to the accelerometer. It can be a comfortable method of monitoring, as an existing wearer of eyeglasses does not need to change their daily habits in order to use it. While not all older adults wear eyeglasses, it may still provide an alternative to other, more obtrusive, methods. A machine learning method was proposed for the automatic detection of epochs of chewing activity from the accelerometer data. The presented methodology was able to achieve an average PPV and sensitivity of 87.8 % and 85.7 %, respectively, on data recorded in a lab environment. This is on par with, and in some cases outperforms, the state of the art evaluated on datasets recorded in a similar environment. When applied to data recorded in a real-life environment, performance drops, as is expected. With data recorded in a nursing home from older adults, an average PPV and sensitivity of 82.7 % and 65.9 % is achieved, respectively. While PPV loses a few percentages, the sensitivity takes the largest hit, indicating that the majority of detections made by the algorithm are correct, but it misses more bite events than in a lab environment.

Furthermore, the thesis introduced a new type of weighing system for the measurement of food bites. A novel smart plate system was introduced as an alternative to table embedded scales. The smart plate is portable, standalone and works with a regular off-the-shelf polymer eating plate, such as the ones commonly used in food delivery services, hospitals or nursing homes. The thesis introduced the mechanical and electrical design methodology of the smart plate prototype with a custom designed embedded system. This prototype was then used to record data in a nursing home and a hospital. The novel aspect of the plate is the ability to localise individual bites on the plate. In combination with a compartmentalised plate, the system can estimate from which compartment the bite was taken, without sensors embedded into the plate itself. All sensor and embedded hardware is contained in the base of the system, the plate remains standard and is removable from the base for easy cleaning and serving of food. With prior knowledge of the food served in each compartment, the localisation can estimate the total amount of calories consumed per food type. The bite localisation algorithm requires an accurate measurement of each individual bite. For this purpose, a novel bite detector based on a supervised learner is proposed. Validated on data recorded in a nursing home and hospital, the bite detector is able to detect bites during meals with an average PPV and

sensitivity of 78 % and 82 %, respectively. Detected bites are then passed to the bite localisation algorithm that assigns a compartment number to each bite. By summing together the bites in each compartment, consumed food can be weighed per compartment. The system was able to measure the weight of food in each compartment with an error ranging from as low as a few percentages to 20 % of the total portion size. In half of the measurements in the dataset, the error was below 10 %. The error highly depends on the eating style of the subject. In some cases, subjects had the tendency to lean on the plate or eat very fast, making it difficult to correctly assign bites to compartments. Even if the compartment detection makes mistakes, the total weight may still be correctly measured. On average, the weight of the total meal could be measured by summing the individual bites with an average error of 9 % of the portion weight.

The final part of the thesis introduced an exploratory study into the quantification of eating behaviour using bite-related parameters as a descriptor of behavioural traits. Eating behaviour is an important aspect that should be included in the overall process of food intake monitoring. The technical realisations presented in the thesis focus more on the quantification of the mechanical action of food intake: measuring activity and the amount of consumed food. These measurements lack contextual data about the *quality* of the eating activity. It has been proposed in the literature that the measurement of individual bites from table embedded scale data may provide insights into the microstructure of eating activity and behaviour. For this purpose, we explored the use of bite-related parameters such as bite length and pickup time to link food intake to certain behavioural traits. In a first part, an attempt is made to correlate bite parameters to the age of the subject, as behavioural traits may change with increasing age. Furthermore, an unlabelled dataset from patients rehabilitating after a stroke was recorded. Patients in this dataset all had an impairment of the upper libs, restricting their ability to fluently consume food. This impairment can be seen as a behavioural trait. Several parameters were proposed that may link this behaviour to bite-related parameters extracted from the smart plate data. The results indicate that such a correlation may exist and may lead the way to further studies with the smart plate for the analysis of behavioural traits.

9.2 Future Work

The work presented in this thesis made several first steps towards the application of monitoring technology for older adults. This may lead the way for further developments and improvements on the presented methodologies. Detecting

food intake during unrestricted eating remains a challenging task, with some challenges remaining unresolved in the thesis. The following sections provide a perspective into potential future improvements.

9.2.1 Accelerometer Based Monitoring

The accelerometer system consists of an off-the-shelf sensor mounted to a glasses frame. The system is not very ergonomic nor aesthetically pleasing. Prior to a real-world application, the system should be miniaturised. Advances in miniature electronics should make it possible to integrate the sensor system into the frame of the glasses. To store and process the data, the same approach as the smart plate can be taken. Data can be temporarily stored on the device and relayed to a different system for processing, for example during charging of the device, to minimise power usage of the device during the day. The algorithm proposed in the thesis should require little computational power and processing on an embedded system should be possible. The most intensive step is the feature extraction, but the chosen features are relatively easy to compute and optimised functions such as the Fast Fourier Transform (FFT) already exist for embedded platforms. Processing is also not required in real-time. Data that was collected during the day could be processed at night when the device is charging.

The accelerometer data recorded and presented in the thesis includes data from wrist worn accelerometers. This data, however, was not used in the study. With a wrist worn accelerometer, different kinds of eating related activities can be measured, such as drinking or preparing food. This would increase complexity of the system, which may not be wanted by older adults, but could still be used in certain situations where a short-term monitoring is required. The data we recorded consists of relatively short measurements of a single meal. These kind of wearable system are expected to be use throughout the day, not only worn during meals. For a complete generalised model, data should be recorded over longer periods of time of at least several days per person. The thesis has made the first step of confirming that the accelerometer can be used to measure chewing activity. In future work, combining the head mounted accelerometer with the body worn sensors could improve the overall detection of eating activity throughout the day.

Furthermore, the performance of the machine learning model is in part limited by the quality of the features. Features used in our study were based on findings from the literature, including features that have shown promising results for the detection of eating activity. There may still be improvements to be made by adding different feature types, or by combining features from

the time and frequency domain. Different detection algorithms can also be looked into to increase performance. Deep learning has emerged as one of the highest performing methods for activity recognition. These techniques, however, typically require very large data sets, whereas our dataset is relatively small.

The current system has only been used to detect chewing activity, but as discussed in Section 5.6, the data may contain additional information that can be used in the study of eating behaviour. In a controlled environment, the chewing activity is visible in the time domain and a link between chewing strength and signal amplitude has been observed. In future work, this could be further looked into as a method to measure chewing strength or ability.

9.2.2 Smart Plate Based Monitoring

The smart plate system currently exists in a prototype phase. While the hardware operates as expected, it is not yet mature enough to be used by individual users or older adults. Currently, a researcher is always present during measurements. In order to facilitate large scale studies, a more user friendly iteration of the system is required. This primarily requires iterations on the hardware. The base platform currently does not have an enclosure and the sensors and PCB are exposed once the plate is removed. Furthermore, the base is currently only compatible with a single type of round polymer plate. This can be further expanded to allow for different kinds of plates to be used. The smart plate methodology may even be applied to eating surfaces. Instead of embedded the hardware underneath a single plate, it can be embedded under an eating surface or integrated into an eating mat. This would allow for regular ceramic plates and bowls to be placed on the mat and position and weight detection could be done in a similar fashion.

Improvements to the bite detection model can further increase the accuracy of the system. Currently, the bite detector performs well with performance scores that improve on the state of the art. There are still some situations during unrestricted eating where the bite detector fails to detect a bite. Misclassification of bites will result in a faulty bite localisation, as an accurate weighing before and after the bite is required for localisation. Currently, localisation is done by taking discrete measurements: one before and one after the bite. The localisation algorithm may be further improved by continuously sampling the weight distribution on the plate and predicting the position of bites using longitudinal data instead of discrete measurements. Similar to the bite detector, features can be extracted from the continuous sampled weight distribution to classify bites in certain compartments, trained with labelled data, instead of the fixed deterministic model that is currently used.

9.2.3 Behavioural Analysis

The study regarding the quantification of behavioural traits from smart plate data is an exploratory study. We showed that it is possible to extract parameters such as bite size, length and pickup time from the bite detection algorithm, but further validation is required. The age-related results are most likely biased due to the limited size of the dataset. In order to get more definitive results, a bigger dataset coupled with ground truth behavioural monitoring methods is required to validate the methodology. Furthermore, behavioural analysis in combination with food intake data can provide new insights for different target groups other than older adults. As shown in our preliminary study, there may be a link between quality of food intake and upper body impairment after stroke.

While the proposed methods may be suited for use by older adults, they are not the only target group that can benefit from the technology. There may be application potential in different situations requiring food intake monitoring. In the treatment and management of obesity, for example, lifestyle intervention and behavioural therapy play an important role. The behavioural therapy consists of daily monitoring of food intake, typically with pen-and-paper food diaries or digital supported applications, which have several shortcomings, as described in the thesis. An automatic monitoring system can be a first step in reducing these shortcomings. Furthermore, the smart plate system and wearable devices can be integrated into the daily routine of patients in a transparent fashion, recording food intake in the background. This can help in weight loss and weight management by increasing awareness of the consumed food through visualisation and personalised feedback derived from the data.

9.2.4 Valorisation Potential

The thesis has been largely written from an academic standpoint. In their current form, the tools and algorithms presented in the thesis are primarily meant to assist researchers or clinical professionals in gaining a better insight into the eating behaviour of older adults. The system can act as an aid in the diagnosis and prevention of malnutrition. When food is served from a central kitchen or food delivery service, as is typically the case in hospitals or nursing homes, the contents and calorie value of the meals are known beforehand. When combined with the compartmentalised smart plate and food types served in different compartments, the calorie content of consumed food can be automatically measured. The application potential is, however, not limited to older adults alone. The tools can be used for other target groups where food intake monitoring is prevalent. This can be a clinical application,

such as the management and treatment of obesity. In this case, the system can act as a therapeutic tool to assist in the management of calorie intake. Finally, the system can be used in research settings as an assessment tool. The impact of various interventions on nutritional intake such as a change in drugs or lifestyle can be measured and included in the research protocol.

Furthermore, there is a potential for commercialisation of the technology into a consumer product. The concept of *quantified-self* has gained popularity in recent years and food intake is already a part of that. Wearable activity trackers on the market typically offer the option of keeping a food diary that can be compared to the daily calorie burn rate. A personal weight goal may be set and data from the wearable is used to create a personalised exercise schedule. The limitations of this approach are similar to those of traditional food intake monitoring systems today. It relies on a manual data entry of the daily consumed food, typically in a smartphone application. In absence of a *smart scale*, body weight needs to be manually recorded as well.

The smart plate system can play a role in simplifying the tracking of daily food intake. Even if food is not served from a central kitchen, but when used at home, food intake can automatically be measured and assigned to a compartment on the plate. The user only needs to enter food types per compartment, which is less labour intensive than the current practice.

Before the technology can be commercialised and turned into a product, however, there are several working points. First is a revision of the hardware. The smart plate is currently a prototype that requires specialised knowledge to use and operate. The base station should be redesigned to make it more robust and better looking. In future iterations it can be made to accept regular eating plates, which would further reduce the adaptation threshold. Given the measurement of three meals per day, the smart plate can currently be used approximately three weeks without charging. A limited amount of effort was placed in the minimisation of energy consumption. There is still room to lower the energy consumption or simply install a bigger battery. This should allow an autonomy well above one month, which is well above the autonomy of other kinds of activity trackers. Second is a further validation of the algorithms in uncontrolled environments. Currently, data was recorded from a small number of users in a select few environments. To ensure the system operates in a variety of settings, further testing and data collection is required. This may also require a revision of the proposed algorithms. Further development of the smart plate can be beneficial for both clinical and consumer applications for the automatic monitoring of food intake.

Bibliography

[1] R.A. DiMaria-Ghalili and E. Amella. Nutrition in older adults: Intervention and assessment can help curb the growing threat of malnutrition. *American Journal of Nursing*, 105:40–50, 2005.

[2] Gianluca Isaia, Simona Mondino, Cristina Germinara, Giorgetta Cappa, Nicoletta Aimonino-Ricauda, Mario Bo, Giovanni Carlo Isaia, Giulia Nobili, and Massimiliano Massaia. Malnutrition in an elderly demented population living at home. *Archives of Gerontology and Geriatrics*, 53(3):249 – 251, 2011.

[3] L.M. Donini, P. Scardella, L. Piombo, B. Neri, R. Asprino, A.R. Proietti, S. Carcaterra, E. Cava, S. Cataldi, D. Cucinotta, G. Di Bella, M. Barbagallo, and A. Morrone. Malnutrition in elderly: Social and economic determintants. *The Journal of Nutrition, Health & Aging*, 17:9–15, 2013.

[4] Nutricia. Results of the NutriAction II study, 2013.

[5] LM Donini, C Savina, M Piredda, D Cucinotta, A Fiorito, EM Inelmen, G Sergi, LJ Dominguez, M Barbagallo, and C Cannella. Senile anorexia in acute-ward and rehabilitation settings. *The Journal of Nutrition Health and Aging*, 12(8):511–517, 2008.

[6] Tanvir Ahmed and Nadim Haboubi. Assessment and management of nutrition in older people and its importance to health. *Clinical interventions in aging*, 5:207, 2010.

[7] Emanuele Cereda, Carlo Pedrolli, Annunciata Zagami, Alfredo Vanotti, Silvano Piffer, Annalisa Opizzi, Mariangela Rondanelli, and Riccardo Caccialanza. Nutritional screening and mortality in newly institutionalised elderly: a comparison between the geriatric nutritional risk index and the mini nutritional assessment. *Clinical Nutrition*, 30(6):793–798, 2011.

[8] Dorothee Volkert, Lioba Pauly, Peter Stehle, and Cornel C Sieber. Prevalence of malnutrition in orally and tube-fed elderly nursing home residents in germany and its relation to health complaints and dietary intake. *Gastroenterology research and practice*, 2011.

[9] H Lochs, C Pichard, and SP Allison. Evidence supports nutritional support. *Clinical Nutrition*, 25(2):177–179, 2006.

[10] L.E. Burke, M. Warziski, T. Starrett, J. Choo, E. Music, S. Sereika, S. Stark, and M.A. Sevick. Self-monitoring dietary intake: Current and future practices. *Journal of Renal Nutrition*, pages 281–290, 2005.

[11] Jin-Ming Wu, Hwan-Jeu Yu, Te-Wei Ho, Xing-Yu Su, Ming-Tsan Lin, and Feipei Lai. Tablet pc-enabled application intervention for patients with gastric cancer undergoing gastrectomy. *Computer methods and programs in biomedicine*, 119(2):101–109, 2015.

[12] Olivier Bouillanne, Gilles Morineau, Claire Dupont, Isabelle Coulombel, Jean-Pierre Vincent, Ioannis Nicolis, Simone Benazeth, Luc Cynober, and Christian Aussel. Geriatric nutritional risk index: a new index for evaluating at-risk elderly medical patients. *The American journal of clinical nutrition*, 82(4):777–783, 2005.

[13] Patricia A Parmelee, Paul D Thuras, Ira R Katz, and M Powell Lawton. Validation of the cumulative illness rating scale in a geriatric residential population. *Journal of the American Geriatrics Society*, 1995.

[14] Oliver Amft. A wearable earpad sensor for chewing monitoring. In *Sensors, 2010 IEEE*, pages 222–227. IEEE, 2010.

[15] Sebastian Passler and Wolf-Joachim Fischer. Food intake monitoring: Automated chew event detection in chewing sounds. *IEEE journal of biomedical and health informatics*, 18(1):278–289, 2014.

[16] Gert Mertes, Hans Hallez, Bart Vanrumste, and Tom Croonenborghs. Detection of chewing motion in the elderly using a glasses mounted accelerometer in a real-life environment. In *Engineering in Medicine and Biology Society (EMBC), 2017 39th Annual International Conference of the IEEE*, pages 4521–4524. IEEE, 2017.

[17] Rui Zhang and Oliver Amft. Monitoring chewing and eating in free-living using smart eyeglasses. *IEEE journal of biomedical and health informatics*, 22(1):23–32, 2018.

[18] E.S. Sazonov and J.M. Fontana. A sensor system for automatic detection of food intake through non-invasive monitoring of chewing. *IEEE Journal of Sensors*, 12:1340–1348, 2012.

[19] J.M. Fontana, M. Farooq, and E. Sazonov. Automatic ingestion monitor: A novel wearable device for monitoring of ingestive behavior. *IEEE Transactions on Biomedical Engineering*, pages 1772–1779, 2014.

[20] Oliver Amft and Gerhard Tröster. Recognition of dietary activity events using on-body sensors. *Artificial Intelligence in Medicine*, 42(2):121–136, 2008.

[21] E. Mendi, O. Ozyavuz, E. Pekesen, and C. Bayrak. Food intake monitoring system for mobile devices. In *IEEE International Workshop on Advances in Sensors and Interfaces*, 2013.

[22] Hongsheng He, Fanyu Kong, and Jindong Tan. Dietcam: Multiview food recognition using a multikernel svm. *IEEE journal of biomedical and health informatics*, 20(3):848–855, 2016.

[23] Rick Weiss, Phyllis J Stumbo, and Ajay Divakaran. Automatic food documentation and volume computation using digital imaging and electronic transmission. *Journal of the American Dietetic Association*, 110(1):42–44, 2010.

[24] Corby K Martin, Sertan Kaya, and Bahadir K Gunturk. Quantification of food intake using food image analysis. In *Engineering in Medicine and Biology Society, 2009. EMBC 2009. Annual International Conference of the IEEE*, pages 6869–6872. IEEE, 2009.

[25] M. Puri, Z. Zhiwei, Y. Qian, A. Divakaran, and H. Sawhney. Recognition and volume estimation of food intake using a mobile device. In *Workshop on Applications of Computer Vision*, 2009.

[26] Anabela Berenguer, Jorge Goncalves, Simo Hosio, Denzil Ferreira, Theodoros Anagnostopoulos, and Vassilis Kostakos. Are smartphones ubiquitous?: An in-depth survey of smartphone adoption by seniors. *IEEE Consumer Electronics Magazine*, 6(1):104–110, 2017.

[27] Pei-Yu Peggy Chi, Jen-Hao Chen, Hao-Hua Chu, and Jin-Ling Lo. Enabling calorie-aware cooking in a smart kitchen. In *International Conference on Persuasive Technology*, pages 116–127. Springer, 2008.

[28] Keng-hao Chang, Shih-yen Liu, Hao-hua Chu, Jane Yung-jen Hsu, Cheryl Chen, Tung-yun Lin, Chieh-yu Chen, and Polly Huang. The diet-aware dining table: Observing dietary behaviors over a tabletop surface. In *International Conference on Pervasive Computing*, pages 366–382. Springer, 2006.

[29] Ryan S Mattfeld, Eric R Muth, and Adam Hoover. Measuring the consumption of individual solid and liquid bites using a table embedded scale during unrestricted eating. *IEEE journal of biomedical and health informatics*, 21(6):1711–1718, 2017.

[30] Bo Zhou, Jingyuan Cheng, Mathias Sundholm, Attila Reiss, Wuhuang Huang, Oliver Amft, and Paul Lukowicz. Smart table surface: A novel approach to pervasive dining monitoring. In *Pervasive Computing and Communications (PerCom), 2015 IEEE International Conference on*, pages 155–162. IEEE, 2015.

[31] Frances E Thompson, Amy F Subar, Catherine M Loria, Jill L Reedy, and Tom Baranowski. Need for technological innovation in dietary assessment. *Journal of the American Dietetic Association*, 110(1):48–51, 2010.

[32] J.M. Fontana and E.S. Sazonov. A robust classification scheme for detection of food intake through non-invasive monitoring of chewing. In *34th Annual International Conference of the IEEE EMBS*, 2012.

[33] P. Lopez-Meyer, S. Schuckers, O. Makeyev, and E. Sazonov. Detection of periods of food intake using support vector machines. In *32nd Annual International Conference of the IEEE EMBS*, 2010.

[34] Edward Sazonov, Stephanie Schuckers, Paulo Lopez-Meyer, Oleksandr Makeyev, Nadezhda Sazonova, Edward L Melanson, and Michael Neuman. Non-invasive monitoring of chewing and swallowing for objective quantification of ingestive behavior. *Physiological measurement*, 29(5):525, 2008.

[35] Sebastian Passler and Wolf-Joachim Fischer. Food intake activity detection using a wearable microphone system. In *Intelligent Environments (IE), 2011 7th International Conference on*, pages 298–301. IEEE, 2011.

[36] Wen Wu and Jie Yang. Fast food recognition from videos of eating for calorie estimation. In *Multimedia and Expo, 2009. ICME 2009. IEEE International Conference on*, pages 1210–1213. IEEE, 2009.

[37] Gianluigi Ciocca, Paolo Napoletano, and Raimondo Schettini. Food recognition: A new dataset, experiments, and results. *IEEE J. Biomedical and Health Informatics*, 21(3):588–598, 2017.

[38] Simon Mezgec and Barbara Koroušić Seljak. Nutrinet: a deep learning food and drink image recognition system for dietary assessment. *Nutrients*, 9(7):657, 2017.

[39] Shulin Yang, Mei Chen, Dean Pomerleau, and Rahul Sukthankar. Food recognition using statistics of pairwise local features. In *Computer Vision and Pattern Recognition (CVPR), 2010 IEEE Conference on*, pages 2249–2256. IEEE, 2010.

[40] Marc Bolanos and Petia Radeva. Simultaneous food localization and recognition. *arXiv preprint arXiv:1604.07953*, 2016.

[41] Wenyan Jia, Yuecheng Li, Ruowei Qu, Thomas Baranowski, Lora E Burke, Hong Zhang, Yicheng Bai, Juliet M Mancino, Guizhi Xu, Zhi-Hong Mao, et al. Automatic food detection in egocentric images using artificial intelligence technology. *Public health nutrition*, pages 1–12, 2018.

[42] Dima Damen, Hazel Doughty, Giovanni Maria Farinella, Sanja Fidler, Antonino Furnari, Evangelos Kazakos, Davide Moltisanti, Jonathan Munro, Toby Perrett, Will Price, and Michael Wray. Scaling egocentric vision: The epic-kitchens dataset. In *European Conference on Computer Vision (ECCV)*, 2018.

[43] P. Pouladzadeh, S Shirmohammadi, and T Arici. Intelligent svm based food intake measurement system. In *IEEE International Conference on Computational Intelligence and Virtual Environments for Measurement Systems and Applications*, 2013.

[44] Fengqing Zhu, Marc Bosch, Insoo Woo, SungYe Kim, Carol J Boushey, David S Ebert, and Edward J Delp. The use of mobile devices in aiding dietary assessment and evaluation. *IEEE journal of selected topics in signal processing*, 4(4):756–766, 2010.

[45] HARRY R Kissileff, GARY Klingsberg, and THEODORE B Van Itallie. Universal eating monitor for continuous recording of solid or liquid consumption in man. *American Journal of Physiology-Regulatory, Integrative and Comparative Physiology*, 238(1):R14–R22, 1980.

[46] Terence M Dovey, David Clark-Carter, Emma J Boyland, and Jason CG Halford. A guide to analysing universal eating monitor data: assessing the impact of different analysis techniques. *Physiology & Behavior*, 96(1):78–84, 2009.

[47] Harry R Kissileff, John Thornton, and Emil Becker. A quadratic equation adequately describes the cumulative food intake curve in man. *Appetite*, 3(3):255–272, 1982.

[48] Esther M Bobroff and Harry R Kissileff. Effects of changes in palatability on food intake and the cumulative food intake curve in man. *Appetite*, 7(1):85–96, 1986.

[49] John D Davis. The microstructure of ingestive behavior. *Annals of the New York Academy of Sciences*, 575(1):106–121, 1989.

[50] HR Kissileff and JL Guss. Microstructure of eating behavior in humans. *Appetite*, 36(1):70, 2001.

[51] Marti A. Hearst, Susan T Dumais, Edgar Osuna, John Platt, and Bernhard Scholkopf. Support vector machines. *IEEE Intelligent Systems and their applications*, 13(4):18–28, 1998.

[52] Ingo Steinwart and Andreas Christmann. *Support vector machines*. Springer Science & Business Media, 2008.

[53] Peter Flach. *Machine learning: the art and science of algorithms that make sense of data*. Cambridge University Press, 2012.

[54] James M Keller, Michael R Gray, and James A Givens. A fuzzy k-nearest neighbor algorithm. *IEEE transactions on systems, man, and cybernetics*, (4):580–585, 1985.

[55] Jerome Friedman, Trevor Hastie, and Robert Tibshirani. *The elements of statistical learning*, volume 1. Springer series in statistics New York, NY, USA:, 2001.

[56] Douglas M Hawkins. The problem of overfitting. *Journal of chemical information and computer sciences*, 44(1):1–12, 2004.

[57] Avrim L Blum and Pat Langley. Selection of relevant features and examples in machine learning. *Artificial intelligence*, 97(1-2):245–271, 1997.

[58] Isabelle Guyon and André Elisseeff. An introduction to variable and feature selection. *Journal of machine learning research*, 3(Mar):1157–1182, 2003.

[59] Andrew Y Ng. Feature selection, l 1 vs. l 2 regularization, and rotational invariance. In *Proceedings of the twenty-first international conference on Machine learning*, page 78. ACM, 2004.

[60] H. Peng, F. Long, and C. Ding. Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27:1226–1238, 2005.

[61] Marie Chan, Eric Campo, Daniel Estève, and Jean-Yves Fourniols. Smart homes-current features and future perspectives. *Maturitas*, 64(2):90–97, 2009.

[62] Veerle Claes, Els Devriendt, Jos Tournoy, and Koen Milisen. Attitudes and perceptions of adults of 60 years and older towards in-home monitoring of the activities of daily living with contactless sensors: an explorative study. *International journal of nursing studies*, 52(1):134–148, 2015.

[63] Youngjung Geum, Hongseok Jeon, and Hakyeon Lee. Developing new smart services using integrated morphological analysis: integration of the market-pull and technology-push approach. *Service Business*, 10(3):531–555, 2016.

[64] Anne-mie Sponselee, Ben Schouten, Don Bouwhuis, and Charles Willems. Smart home technology for the elderly: Perceptions of multidisciplinary stakeholders. In *European Conference on Ambient Intelligence*, pages 314–326. Springer, 2007.

[65] Greet Baldewijns, Patricia Sabbe, Kristien Rombouts, Kurt Peeters, An Mondelaers, Jessica Hekking, Tom Croonenborghs, and Bart Vanrumste. Establishing a collaboration between care providers and engineers. In *Raising Awareness for the Societal and Environmental Role of Engineering and (Re) Training Engineers for Participatory Design (Engineering4Society), 2015 Conference on*, pages 17–20. IEEE, 2015.

[66] LH Segura Anaya, Abeer Alsadoon, Nectar Costadopoulos, and PWC Prasad. Ethical implications of user perceptions of wearable devices. *Science and engineering ethics*, pages 1–28, 2018.

[67] Yu-Min Fang and Chien-Cheng Chang. Users' psychological perception and perceived readability of wearable devices for elderly people. *Behaviour & Information Technology*, 35(3):225–232, 2016.

[68] Robert Steele, Amanda Lo, Chris Secombe, and Yuk Kuen Wong. Elderly persons' perception and acceptance of using wireless sensor networks to assist healthcare. *International journal of medical informatics*, 78(12):788–801, 2009.

[69] Cherrylyn Buenaflor and Hee-Cheol Kim. Wearable computers in human perspective: the decision process of innovation acceptance with user issues and concerns. *International Journal of Emerging Technology and Advanced Engineering*, 2(11):573–580, 2012.

[70] Ronald E McGaughey, Steven M Zeltmann, and Mark E McMurtrey. Motivations and obstacles to smartphone use by the elderly: developing a research framework. *International Journal of Electronic Finance*, 7(3-4):177–195, 2013.

[71] Chenkai Wu, Ellen Smit, Qian-Li Xue, and Michelle C Odden. Prevalence and correlates of frailty among community-dwelling chinese older adults: the china health and retirement longitudinal study. *The Journals of Gerontology: Series A*, 73(1):102–108, 2017.

[72] Xue-Qiang Wang and Pei-Jie Chen. Population ageing challenges health care in china. *The Lancet*, 383(9920):870, 2014.

[73] United Nations Department of Economic and Social Affairs: Population Division. *World Population Prospects: The 2015 Revision*. United Nations Publications, 2015.

[74] World Health Organization. China country assessment report on ageing and health. 2015.

[75] JunDe Zhou, Miao Wang, HaiKuan Wang, and Qiang Chi. Comparison of two nutrition assessment tools in surgical elderly inpatients in northern china. *Nutrition journal*, 14(1):68, 2015.

[76] NC Shum, WWH Hui, FCS Chu, J Chai, and TW Chow. Prevalence of malnutrition and risk factors in geriatric patients of a convalescent and rehabilitation hospital. *Hong Kong Medical Journal*, 2005.

[77] Luxi Ji, Hongdao Meng, and Birong Dong. Factors associated with poor nutritional status among the oldest-old. *Clinical nutrition*, 31(6):922–926, 2012.

[78] Ling Bao and Stephen S Intille. Activity recognition from user-annotated acceleration data. In *International Conference on Pervasive Computing*, pages 1–17. Springer, 2004.

[79] Adil Mehmood Khan, Young-Koo Lee, Sungyoung Y Lee, and Tae-Seong Kim. A triaxial accelerometer-based physical-activity recognition via augmented-signal features and a hierarchical recognizer. *IEEE transactions on information technology in biomedicine*, 14(5):1166–1172, 2010.

[80] Nishkam Ravi, Nikhil Dandekar, Preetham Mysore, and Michael L Littman. Activity recognition from accelerometer data. In *Aaai*, volume 5, pages 1541–1546, 2005.

[81] Jamie A Ward, Paul Lukowicz, Gerhard Troster, and Thad E Starner. Activity recognition of assembly tasks using body-worn microphones and accelerometers. *IEEE transactions on pattern analysis and machine intelligence*, 28(10):1553–1567, 2006.

[82] Uwe Maurer, Asim Smailagic, Daniel P Siewiorek, and Michael Deisher. Activity recognition and monitoring using multiple sensors on different body positions. In *Wearable and Implantable Body Sensor Networks, 2006. BSN 2006. International Workshop on*, pages 4–pp. IEEE, 2006.

[83] Adil Mehmood Khan, Y-K Lee, SY Lee, and T-S Kim. Human activity recognition via an accelerometer-enabled-smartphone using kernel discriminant analysis. In *Future Information Technology (FutureTech), 2010 5th International Conference on*, pages 1–6. IEEE, 2010.

[84] Oscar D Lara, Miguel A Labrador, et al. A survey on human activity recognition using wearable sensors. *IEEE Communications Surveys and Tutorials*, 15(3):1192–1209, 2013.

[85] Fabio Bagala, Clemens Becker, Angelo Cappello, Lorenzo Chiari, Kamiar Aminian, Jeffrey M Hausdorff, Wiebren Zijlstra, and Jochen Klenk. Evaluation of accelerometer-based fall detection algorithms on real-world falls. *PloS one*, 7(5):e37062, 2012.

[86] Eduardo Casilari, José-Antonio Santoyo-Ramón, and José-Manuel Cano-García. Analysis of public datasets for wearable fall detection systems. *Sensors*, 17(7):1513, 2017.

[87] A Van der Bilt, L Engelen, LJ Pereira, HW Van der Glas, and JH Abbink. Oral physiology and mastication. *Physiology & behavior*, 89(1):22–27, 2006.

[88] Andries Van Der Bilt, RAFAEL DE LIZ POCZTARUK, and Jan Hendrik Abbink. Skull vibration during chewing of crispy food. *Journal of texture studies*, 41(6):774–788, 2010.

[89] Shuangquan Wang, Gang Zhou, Yongsen Ma, Lisha Hu, Zhenyu Chen, Yiqiang Chen, Hongyang Zhao, and Woosub Jung. Eating detection and chews counting through sensing mastication muscle contraction. *Smart Health*, 2018.

[90] OpenStax College. *Anatomy and physiology*. Rice University, 2015.

[91] Refet Firat Yazicioğlu. Surface micromachined capacitive accelerometers using mems technology. Master's thesis, Middle East Technical University, 2003.

[92] RM Boysel and LJ Ross. Development of a single-mass five-axis mems motion sensor. In *Advanced Microsystems for Automotive Applications 2009*, pages 333–353. Springer, 2009.

[93] Peter A Flach. The geometry of roc space: understanding machine learning metrics through roc isometrics. In *Proceedings of the 20th International Conference on Machine Learning (ICML-03)*, pages 194–201, 2003.

[94] David A Cieslak and Nitesh V Chawla. Learning decision trees for unbalanced data. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*, pages 241–256. Springer, 2008.

[95] A Sheiham and J Steele. Does the condition of the mouth and teeth affect the ability to eat certain foods, nutrient and dietary intake and nutritional status amongst older people? *Public health nutrition*, 4(3):797–803, 2001.

[96] Caroline C Horwath. Chewing difficulty and dietary intake in the elderly. *Journal of Nutrition for the Elderly*, 9(2):17–24, 1990.

[97] H Miura, M Kariyasu, K Yamasaki, and Y Arai. Evaluation of chewing and swallowing disorders among frail community-dwelling elderly individuals. *Journal of oral rehabilitation*, 34(6):422–427, 2007.

[98] Y Takata, T Ansai, S Awano, M Fukuhara, K Sonoki, M Wakisaka, K Fujisawa, S Akifusa, and T Takehara. Chewing ability and quality of life in an 80-year-old population. *Journal of oral rehabilitation*, 33(5):330–334, 2006.

[99] KAORU KOHYAMA and LAURENCE MIOCHE. Chewing behavior observed at different stages of mastication for six foods, studied by electromyography and jaw kinematics in young and elderly subjects. *Journal of Texture Studies*, 35(4):395–414, 2004.

[100] Jeffrey M Brunstrom and Peter J Rogers. How many calories are on our plate? expected fullness, not liking, determines meal-size selection. *Obesity*, 17(10):1884–1890, 2009.

[101] SV Gupta. Strain gauge load cells. In *Mass Metrology*, pages 89–119. Springer, 2012.

[102] Nicolien Zijlstra, René de Wijk, Monica Mars, Annette Stafleu, and Cees de Graaf. Effect of bite size and oral processing time of a semisolid food on satiation–. *The American journal of clinical nutrition*, 90(2):269–275, 2009.

[103] Sandra F Simmons and David Reuben. Nutritional intake monitoring for nursing home residents: a comparison of staff documentation, direct observation, and photography methods. *Journal of the American Geriatrics Society*, 48(2):209–213, 2000.

[104] Jonathan van't Riet, Siet J Sijtsema, Hans Dagevos, and Gert-Jan De Bruijn. The importance of habits in eating behaviour. an overview and recommendations for future research. *Appetite*, 57(3):585–596, 2011.

[105] AP Hearty, SN McCarthy, JM Kearney, and MJ Gibney. Relationship between attitudes towards healthy eating and dietary behaviour, lifestyle and demographic factors in a representative sample of irish adults. *Appetite*, 48(1):1–11, 2007.

[106] Jane Wardle, Carol Ann Guthrie, Saskia Sanderson, and Lorna Rapoport. Development of the children's eating behaviour questionnaire. *The Journal of Child Psychology and Psychiatry and Allied Disciplines*, 42(7):963–970, 2001.

[107] AJ Hill, PJ Rogers, and JE Blundell. Techniques for the experimental measurement of human eating behaviour and food intake: a practical guide. *International journal of obesity and related metabolic disorders: journal of the International Association for the Study of Obesity*, 19(6):361–375, 1995.

[108] Wagner Marcenes, Jimmy George Steele, Aubrey Sheiham, and Angus Willian Gilmour Walls. The relationship between dental status, food selection, nutrient intake, nutritional status, and body mass index in older people. *Cadernos de saude publica*, 19:809–815, 2003.

[109] Axel R Fugl-Meyer. Post-stroke hemiplegia assessment of physical properties. *Scand J Rehabil Med Suppl*, 7:85–93, 1980.

[110] David J Gladstone, Cynthia J Danells, and Sandra E Black. The fugl-meyer assessment of motor recovery after stroke: a critical review of its measurement properties. *Neurorehabilitation and neural repair*, 16(3):232–240, 2002.

List of Publications

Papers submitted and under review

Gert Mertes, Li Ding, Wei Chen, Hans Hallez, Jie Jia and Bart Vanrumste Measuring and localizing individual bites using a sensor augmented plate during unrestricted eating for the aging population. *IEEE Journal of Biomedical and Health Informatics.* (in revision).

Gert Mertes, Li Ding, Wei Chen, Hans Hallez, Jie Jia and Bart Vanrumste Quantifying eating behavior with a smart plate in patients with arm impairment after stroke. *IEEE International Conference on Biomedical and Health Informatics (BHI 2019).* (in review).

Published papers

Gert Mertes, Glenn Christiaensen, Hans Hallez, Sammy Verslype, Wei Chen, and Bart Vanrumste. Measuring weight and location of individual bites using a sensor augmented smart plate. In *Proc. of the 40th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC).* IEEE, 2018.

Gert Mertes, Hans Hallez, Bart Vanrumste, and Tom Croonenborghs. Detection of chewing motion in the elderly using a glasses mounted accelerometer in a real-life environment. In *Proc. of the 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 4521–4524. IEEE, 2017.

Greet Baldewijns, Glen Debard, Gert Mertes, Tom Croonenborghs, and Bart Vanrumste. Improving the accuracy of existing camera based fall detection algorithms through late fusion. In *Proc. of the 39th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pages 2667–2671. IEEE, 2017.

Gert Mertes, Tom Croonenborghs, Bart Vanrumste, and Hans Hallez. Towards detection of chewing motion in the elderly using a glasses mounted accelerometer. In *Proc. of the International Conference on Biomedical and Health Informatics (BHI)*. IEEE, 2017.

Greet Baldewijns, Glen Debard, Gert Mertes, Bart Vanrumste, and Tom Croonenborghs. Bridging the gap between real-life data and simulated data by providing a highly realistic fall dataset for evaluating camera-based fall detection algorithms. *Healthcare technology letters*, 3(1):6–11, 2016.

Gert Mertes, Hans Hallez, Tom Croonenborghs, and Bart Vanrumste. Detection of chewing motion using a glasses mounted accelerometer towards monitoring of food intake events in the elderly. In *Proc. of the International Conference on Biomedical and Health Informatics (ICBHI)*, 2015.

Gert Mertes, Greet Baldewijns, Tom Croonenborghs, and Bart Vanrumste. Using the Nintendo Wii balance board for automatic fall risk classification. In *Proc. of the 5th Dutch conference on biomedical engineering*, p 165. 2015

Gert Mertes, Greet Baldewijns, Pieter-Jan Dingenen, Tom Croonenborghs, and Bart Vanrumste. Automatic fall risk estimation using the nintendo wii balance board. In *Proc. of the International Conference on Health Informatics*, pages 75–81. SCITEPRESS, 2015.

Science outreach

Gert Mertes. Healthy Ageing: Technology for independence. Research In Flanders Thematic Paper, pages 15-16, 2016.

Greet Baldewijns, Gert Mertes, Hans Hallez, Tom Croonenborghs, and Bart Vanrumste. Technologie op rust : het rusthuis van de toekomst. Clickx Magazine, pages 24-27, 2016.

Greet Baldewijns, Gert Mertes, Hans Hallez, Tom Croonenborghs, and Bart Vanrumste. Ingenieurs in het rusthuis. Capuskrant, 27 (1), pages 2-3, Jan 2016.

Greet Baldewijns and Gert Mertes. Interview met Greet Baldewijns en Gert Mertes: Ingenieurs nemen intrek in rusthuis. Het Laatste Nieuws (Leuven/Brabant), Nov 2015.

Greet Baldewijns, Gert Mertes, Hans Hallez, Tom Croonenborghs, and Bart Vanrumste. Interview met Greet Baldewijns en Gert Mertes: Ingenieurs actief in woonzorgcentrum. Nieuwsblad (Leuven-Hageland), Nov 2015.

Greet Baldewijns, Gert Mertes, Hans Hallez, Tom Croonenborghs, and Bart Vanrumste. Interview: Ingenieurs@wzc. ROB TV, Nov 2015.

Gert Mertes, Hans Hallez, Tom Croonenborghs, and Bart Vanrumste. Bril bestrijdt ondervoeding bij ouderen. Het Belang Van Limburg, Nov 2015



FACULTY OF ENGINEERING TECHNOLOGY
DEPARTMENT OF ELECTRICAL ENGINEERING
E-MEDIA RESEARCH LAB
Andreas Vesaliusstraat 13
B-3000 Leuven
http://www.kuleuven.be/emedia