Nurturing a Community of Digital Leaders

Faculty of ICT Publication 2020

L-Università ta' Malta

Excellence and Trust in Artificial Intelligence

The EU's approach to Artificial Intelligence (AI), based on trust and excellence, will give citizens the confidence to embrace these technologies while encouraging businesses to develop them.

CITIZENS

Better healthcare, safer and cleaner transport and improved public services.

BUSINESSES

Innovative products and services, for example in energy, security, healthcare; higher productivity and more efficient manufacturing.

GOVERNMENTS

Cheaper and more sustainable services such as trasnport, energy and waste management.

Source: European Commission

What is a high-risk AI application?

- When it concerns a critical use in a critical sector

- For example: medical equipment, automated driving, decisions on social security payments;
- Some uses are critical in all sectors, for example use of AI in recruitment processes.

How to enforce trusthworthy AI in practise?

- High-risk AI will be subject to strict rules (compliance tests, controls, sanctions);
- Other AI applications can use voluntary labelling.

Nurturing a community of digital leaders

THE FACULTY OF ICT 2020 PUBLICATION

Dr Conrad Attard and Mr Iggy Fenech

Welcome to The Faculty of ICT 2020 Publication, which complements this year's theme of nurturing a community of digital leaders. And who better to represent these digital leaders than those researchers currently reading for a Master's degree or for a PhD? After all, there is no one better to inspire the next generation of students; who will be paving the way for how we will live tomorrow.

Yet, whether you are a student, a player in industry, or a government representative, there's a lot for you to discover in this publication. Among the research projects you can read about here are those related to archaeology, healthcare, the internet of things (IoT), pervasive computing, FinTech, air quality, mapmaking, and aviation.

These studies are but a drop in the ocean when compared to all the on-going research currently taking place within the Faculty here at the University of Malta. Sadly, we couldn't shine a light on each and every one of them, but these articles, fashioned from interviews done directly with the researchers themselves, should help you get a good idea of the scope of ICT and just how varied research in this area can be.

The researchers mentioned in this publication, meanwhile, will be joined by others in the years to come. Indeed, today's undergraduate students are among the many contributing to the Faculty's success, which is why we have included around 65 abstracts of undergraduate projects submitted this year. Among the many areas that will be worked on are those on Networks & Telecommunications, IoT, Audio, Speech and Language Technology, Software Engineering & Web Applications, Data Science, Blockchain & Fintech, Digital Health, Deep Learning, and Testing and Verification.

Before you start reading, however, it must be said that it's quite a feat that we have managed to put such a publication together in a year like this one. The current pandemic has threatened to spell the end of numerous research projects, but with determination, the Faculty adjusted its curriculum, moved its lessons online, gave students the opportunity to work in digital labs from the safety of their own homes, and extended deadlines where needed.

Life — although different than it was just seven months ago — must continue, and so does research to ensure we live in a better world in the years to come. Indeed, the ultimate aim of this publication is to showcase just how important ICT is to the modern world and how the research conducted through it forms part of the building blocks of tomorrow's societies. And, while all research in ICT revolves around the concepts of artificial intelligence, IoT, and emerging technologies, among others, it's also important to keep in mind that, at its core, ICT is about human beings looking to solve problems, find solutions, and lead better lives.

With that, we will leave you to go through the articles we have put together for you. Enjoy reading the stories and marveling at the research, but also remember that it could be you we feature in our future editions. After all, all these researchers were once prospective students, too.

Happy reading!

A word from the Dean

The Faculty of ICT 2020 Publication also marks the start of another academic year here at the Faculty of Information & Communications Technology (ICT) within the University of Malta. Yet, as a new wave of students joins those returning to the degree courses we offer, things are slightly different. With COVID-19 still making the rounds, students' time on campus has been reduced, but our determination to teach the next generation hasn't.

Although we are a young faculty — formed just 13 years ago in 2007 — I am proud of how responsive we have been to the new reality, and of how my staff has met this challenge head on, both by migrating all lectures online, as well as by making virtual labs available to students. Then again, all this forms part of how important ICT can be, and is, to everyday life.

As time progresses, more and more people are coming to realise the value digital economy and digital software can bring to their lives, particularly during this pandemic, which has accelerated the rise of the digital citizen. But the digital world requires skilled and knowledgeable professionals that are quick on their feet in a fast-changing world. Making sure we prepare today's students to be the professionals our society will need tomorrow is our biggest task, but it's one we are proud to execute.

Across our five departments – namely those of Artificial Intelligence, Communications & Computer Engineering, Computer Information Systems, Computer Science, and Microelectronics & Nanoelectronics – our students and lecturers are slowly but surely helping the world progress. This can be seen in this publication, which shows how many students at Master's and Doctoral levels are working on incredible projects together with other Faculties, foreign universities, as well as other entities and private businesses, both locally and abroad. Meanwhile, a group of Bachelor's students is well on its way to joining them.

"As time progresses, more and more people are coming to realise the value digital economy and digital software can bring to their lives"

Indeed, part of what has made the Faculty of ICT successful is our constant commitment to research, and the fields of study include Health, Archaeology, Aviation and Accessibility, amongst many others. This, on top of helping the Faculty thrive, also gives our students incredible experiences locally and at an EU level. Suffice to mention our Faculty collaborates closely with CERN, the European Organization for Nuclear Research, and has also worked with Dassault Aviation, which gave our students the honour of being amongst the first in the world to model radio propagation of 3G and Wi-Fi signals on aircraft. These experiences shape their studies and are testament to the fact that opportunities are available to those who want them.

Moreover, our work also involves working with policymakers for anything related to ICT in Malta. It's a great responsibility that has a direct impact on people's lives, but it also shows how important and respected the work and research done by our academics and students is.

With that, I will leave you to pore over this new issue. But I must make one concluding remark: all that is mentioned above makes us hopeful for the future and we look forward to sharing these experiences with more students and the general public — after all, the point of ICT has always been to support humanity and to develop solutions that help it advance.

Dr Ing Carl James Debono
Dean of the Faculty of ICT

Kickstart Your Career

Are you ready for your next adventure in FinTech?

We offer you exciting opportunities to accelerate your personal and professional growth in a global company.

Become part of RS2 and work in an international team to forge cutting-edge technologies and payment solutions.

Join us to shape the future of payments

See our latest job offers at **RS2.COM**

recruitment@rs2.com | +356 2134 5857

A CAREER ON THE CUTTING EDGE

A career in STEM positions you at the fore-front of the next ground-breaking technology. Working with leading-edge technology fields will give you the freedom and flexibility in your job to clear a path for new ideas and innovation.

By choosing to pursue a major in a STEM field, you can expect plenty of job availability, a high salary, flexibility, job satisfaction and the opportunity to make an important impact on society.

STEM education creates critical thinkers and enables the next generation of innovators. Exploring this field can result in a successful and fulfilling career—and life.

1 Industry Growth

STEM graduates are in extremely high demand, meaning you are unlikely to experience a shortage of available jobs following graduation.

The fact that STEM salaries continue to rise is a testament to the demand for talented STEM workers. Not only will you earn an impressive salary, but you're also likely to have less difficulty finding employment.

If you want a career that is high-paying, fast-growing, and satisfying, a STEM major is a very solid choice.

2 Post-Graduation Job Security

There is no doubt that students studying STEM have more opportunities to play the field when looking for jobs. The world is your oyster when you have a STEM degree. STEM employers also have a history of providing employees with perks and competitive pay in return for expertise.

There are 3% of all the workers in the EU employed in the sector of advanced technology

Source: Eurostat

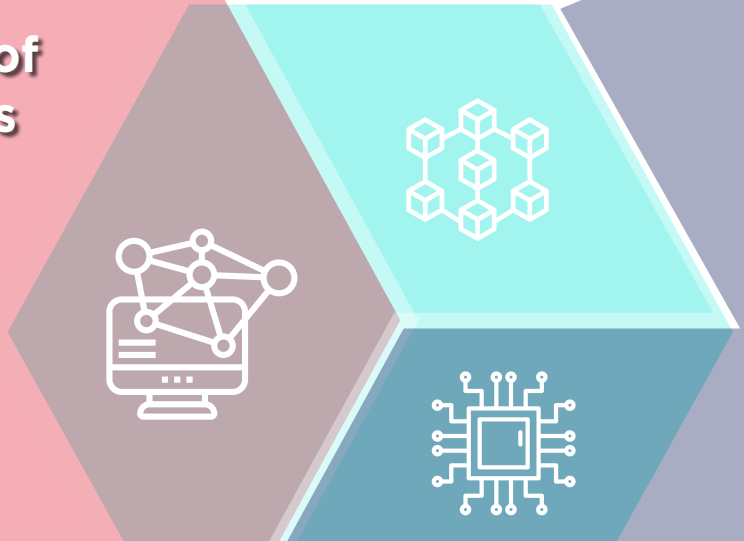

In Malta, this sector has also seen the quickest rate of growth among all EU member states at almost 4% per annum

> **STEM workers are on the cutting edge of innovation in fields like computer technology, artificial intelligence, medicine, engineering, design, robotics and more.**
>
> Dana Farrugia, CEO - Tech.mt

3 Flexibility

While some degrees limit students to only a few jobs, specialising in STEM majors typically have much more flexibility, because your skills are cross-marketable.

These skills are in demand for a wide variety of jobs. It is much easier for STEM majors to switch careers if needed or desired.

STEM education will also provide you with the necessary skills that will enable you to work remotely for a company, become a freelancer or work solely on your own projects.

4 Learn Transferable Skills

You will learn a range of transferable skills that can be used in just about any occupation or industry. You will also boast high-level communication, interpersonal, critical thinking, problem-solving and adaptability skills — attributes that are increasingly important to employers in today's competitive job market. These are known as the new soft skills!

5 You can feed your curiosity

STEM workers are on the cutting edge of innovation in fields like computer technology, artificial intelligence, medicine, engineering, design, robotics and more.

The ability to discover, innovate, and ultimately impact society is one reason that STEM majors report such high levels of job satisfaction.

Cigroup

YOUR PURPOSE IS YOUR BIGGEST ASSET

www.cigroup.mt

CDENewsAgency CiNext CiConsulta #ML InsightsInstitute boldidentities

exigy

LOOKING FOR INSPIRED MINDS

We are a team of industry-leading experts who have a passion for technology. Do you have what it takes to join our team?

Apply now at www.exigy.com/join-our-team

GET IN TOUCH!

(+356) 2011 2000

recruitment@exigy.com

- facebook.com/exigy
- linkedin.com/company/exigy
- www.exigy.com

Nurturing a Community of Digital Leaders

Faculty of ICT Publication 2020

L-Università ta' Malta

Front cover and booklet designed by
Yellow Door Media

Printing by
www.deltamedia.services

Editorial board
Mr Iggy Fenech & Dr Conrad Attard

Abstracts (main text only) proofread by:
Colette Grech @restylelinguistic

Review of Abstracts
Dr Chris Porter & Dr Chris Columbo

Lead Administrator of Publication
Ms Anabel Decesare

Administration of Publication
Ms Michelle Agius
Mr Rene Barun
Mr Anthony Buhagiar
Ms Therese Caruana
Mr Kevin Cortis

Photography
Mr James Moffett
Ms Sarah Zammit

#FICT20

Acknowledgements

The Faculty of Information and Communication Technology gratefully acknowledges the following firms and organisations for supporting this year's Faculty of ICT Publication 2020:

Gold Sponsors

Main Sponsor of the event

Event Sponsors

Exhibition Map

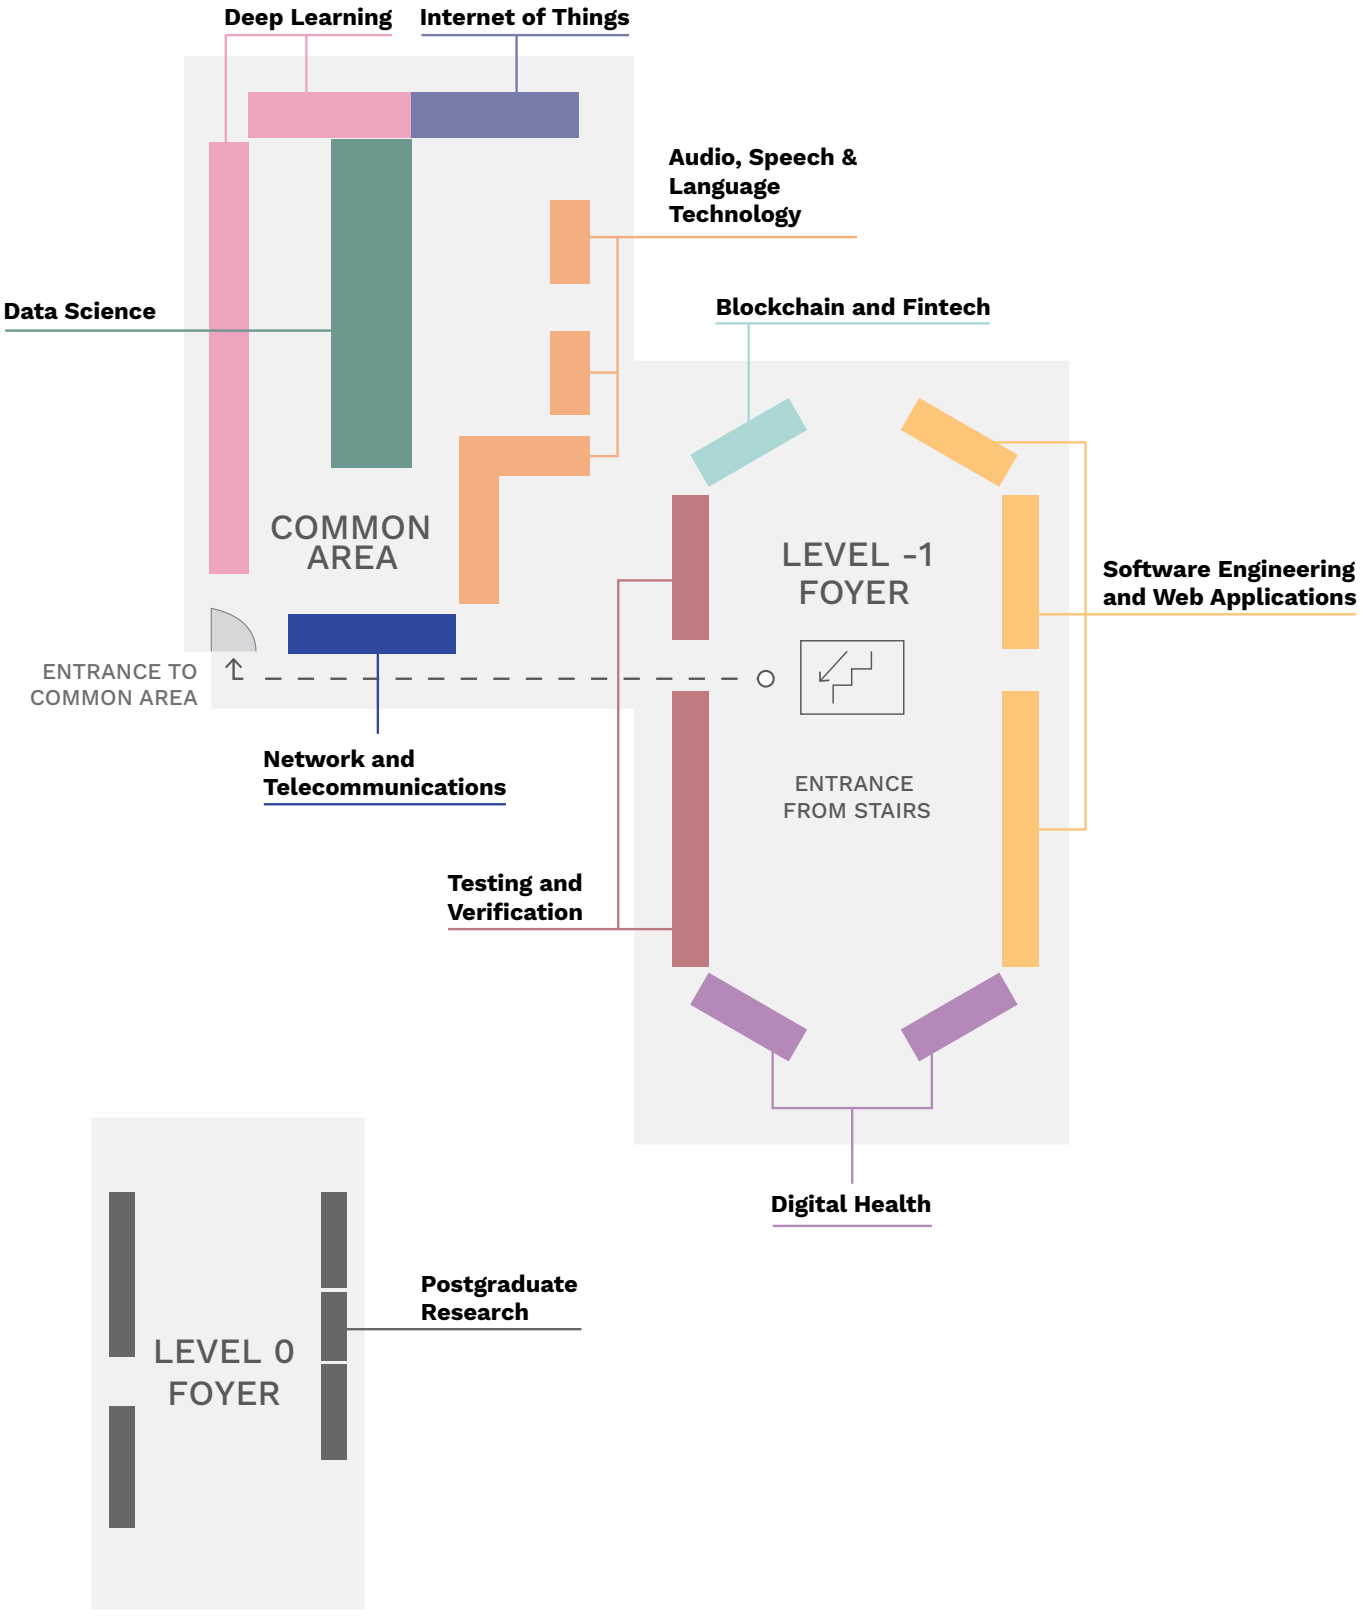

Contents

Blockchain & Fintech

Testing & Verification

Deep Learning

Data Science

Software Engineering & Web Applications

Internet of Things

Audio, Speech & Language Technology

Digital Health

Network & Telecommunications

Members of Staff FACULTY OF ICT

DEPARTMENT OF COMMUNICATIONS & COMPUTER ENGINEERING

PROFESSOR

Professor Inġ. Carl J. Debono, B.Eng.(Hons.), Ph.D.(Pavia), M.I.E.E.E., M.I.E.E.
(Dean of Faculty)

ASSOCIATE PROFESSORS

Professor Johann A. Briffa, B.Eng. (Hons)(Melit.), M.Phil.(Melit.), Ph.D.(Oakland), (Head of Department)
Professor Inġ. Victor Buttigieg, B.Elec.Eng.(Hons.), M.Sc. (Manc.), Ph.D.(Manc.), M.I.E.E.E.
Professor Inġ. Adrian Muscat, B.Eng. (Hons.), M.Sc. (Brad.), Ph.D.(Lond.), M.I.E.E.E.
Professor Inġ. Saviour Zammit, B.Elec.Eng.(Hons.), M.Sc. (Aston), Ph.D.(Aston), M.I.E.E.E.
(Pro-Rector for Research and Innovation)

SENIOR LECTURERS

Dr Inġ. Reuben A. Farrugia, B.Eng.(Hons.), Ph.D., M.I.E.E.E.
Dr Inġ. Trevor Spiteri, B.Eng.(Hons.), M.Sc., Ph.D.(Bris.), M.I.E.E.E., M.I.E.T.

LECTURER

Dr Inġ. Gianluca Valentino, B.Sc.(Hons.)(Melit.), Ph.D. (Melit.), M.I.E.E.E.

AFFILIATE PROFESSOR

Dr Hector Fenech, B.Sc. (Eng.) Hons., M.E.E. (P.I.I.), Ph.D. (Bradford), Fellow A.I.A.A., F.I.E.E.E., F.I.E.T., Eur. Eng.

ASSISTANT LECTURER

Inġ. Etienne-Victor Depasquale, B.Elec.Eng.(Hons.), M.Sc.(Eng.), M.I.E.E.E.

VISITING ASSISTANT LECTURERS

Inġ. Brian E. Cauchi, B.Sc.IT (Hons.), M.Sc.(ICT), M.Ent.
Inġ. Antoine Sciberras, B.Eng.(Hons.)(Melit.), PG.Dip.Eng.Mangt.(Brunel), M.ent (Melit.)
Inġ. Leslie Spiteri, B.Elec.Eng.(Hons.), M.Sc., M.I.E.E.E.
Inġ. Martin Zammit, B.Elec. Eng. (Hons.)

RESEARCH SUPPORT OFFICERS

Ms Leanne Attard, B.Eng.(Hons.), M.Sc. (Research Support Officer)
Dr Arkadiusz Gorzawski, Ph.D. (Research Support Officer I)
Dr Christian Galea, Ph.D (Merit), M.Sc (Melit.), B.Sc. (Hons.) ICT (CCE), MIEEE (Research Support Officer III)
Dr Frederik Van Der Vekan, Ph.D (Research Support Officer)
Dr David Lloyd, MSci (Lond.), DPhil (Oxon.) (Research Support Officer III)
Mr Matthew Sacco, B.Sc (Hons) (Research Support Officer I)
Mr Leander Grech, B.Sc (Hons) (Research Support Officer)
Mr Simon Hirlander, B.Sc (Hons), M.Sc (TU Wien) (Research Support Officer)
Dr Ewan Hamish Maclean, M.Phys, D.Phil (Research Support Officer)
Dr Mang Chen (Research Support Officer III)
Dr Marc Tanti (Research Support Officer III)
Mr Matthew Aquilina (Research Support Officer II)
Mr Riccardo Illan Fiastre (Research Support Officer)

ADMINISTRATIVE & TECHNICAL STAFF

Mr Mark Anthony Xuereb, (Administrator I)
Mr Albert Sacco, (Senior Laboratory Officer)
Inġ. Maria Abela-Scicluna, B.Eng.(Hons.)(Melit.), M.Sc. ICT (Melit.) (Systems Engineer)

Research Areas:

Computer Networks and Telecommunications

- Error Correction Codes
- Multimedia Communications
- Multi-view video coding and transmission
- Video Coding

Signal processing and Pattern Recognition

- Computer Vision
- Image Processing
- Light Field Imaging
- Volumetric Image Segmentation
- Medical Image Processing and Coding
- Earth Observation
- Vision and Language in Robotics
- Visual Relation Detection
- Semantic Depth Estimation
- Simulation in Machine Learning
- Machine Learning and Applications

Computer Systems Engineering

- Data Acquisition and Control Systems for Particle Accelerators and Detectors
- Digital Games Platforms
- Demand Responsive Transport Systems
- Implementation on Massively Parallel Systems (e.g. GPUs)
- Reconfigurable Hardware

DEPARTMENT OF COMPUTER SCIENCE

PROFESSOR
Professor Gordon J. Pace, B.Sc., M.Sc. (Oxon.), D.Phil. (Oxon.)

ASSOCIATE PROFESSOR
Professor Adrian Francalanza, B.Sc.I.T. (Hons.), M.Sc., D.Phil.(Sussex)
Professor Kevin Vella, B.Sc., Ph.D. (Kent)

SENIOR LECTURERS
Dr Mark Micallef, B.Sc.(Hons.), Ph.D. (Melit.), M.B.A.(Melit.) (Head of Department)
Dr Mark J. Vella, B.Sc.I.T.(Hons.), M.Sc. Ph.D. (Strath.)
Dr Joshua Ellul, B.Sc.I.T. (Hons.), M.Sc. (Kent) , Ph.D. (Soton)
Dr Christian Colombo, B.Sc.I.T. (Hons.), M.Sc. Ph.D. (Melit.)
Dr Keith Bugeja, B.A.(Hons), M.IT, Ph.D.(Warw.)

LECTURERS
Dr Sandro Spina, B.Sc.I.T.(Hons), M.Sc. (Melit), Ph.D.(Warw.)

AFFILIATE LECTURER
Dr Neville Grech, B.Sc.(Hons),M.Sc.(S'ton),Ph.D.(S'ton)

RESEARCH SUPPORT OFFICERS
Ms Caroline Caruana B.Sc.(Melit.), M.Sc.(Melit.) (Research Support Officer I)
Mr Mark Charles Magro, B.Sc.(Melit.),M.Sc.(Melit.) (Research Support Officer II)
Mr Adrian De Barro, B.Sc.ICT(Hons)(Melit.),M.Sc.(Melit.) (Research Support Officer II)
Mr Kevin Napoli, B.Sc.ICT(Hons)(Melit.),M.Sc.(Melit.) (Research Support Officer II)
Ms Jennifer Bellizzi, B.Sc.ICT(Hons)(Melit.), M.Sc.(Birmingham) (Research Support Officer II)
Mr Yonas Leguessei B.Sc.(Melit.), M.Sc.(Melit.) (Research Support Officer II)
Mr Robert Abela, B.Sc.(Hons), M.Sc.(Melit.) (Research Support Officer II)

ADMINISTRATIVE & TECHNICAL STAFF
Mr Kevin Cortis, B.A.(Hons) Graphic Design & Interactive Media (Administrator II))

Research Areas:

- Concurrency
- Compilers
- Distributed Systems and Distributed Ledger Technologies
- Model Checking and Hardware/Software Verification
- Operating Systems
- Semantics of Programming Languages
- High Performance Computing and Grid Computing
- Runtime Verification
- Software Development Process Improvement and Agile Processes
- Software Engineering
- Software Testing
- Security

DEPARTMENT OF MICROELECTRONICS AND NANOELECTRONICS

PROFESSOR

Professor Inġ. Joseph Micallef, B.Sc.(Eng.)(Hons.),M.Sc.(Sur.),Ph.D.(Sur.), M.I.E.E.E.
Professor Ivan Grech, B.Eng.(Hons.),M.Sc.,Ph.D.(Sur.),M.I.E.E.E.

ASSOCIATE PROFESSORS

Professor Inġ. Edward Gatt, B.Eng.(Hons.),M.Phil.,Ph.D.(Sur.),M.I.E.E.E.

SENIOR LECTURERS

Dr Inġ. Owen Casha, B. Eng.(Hons.) (Melit.),Ph.D. (Melit.), M.I.E.E.E. (Head of Department)
Dr Inġ. Nicholas Sammut, B.Eng.(Hons.) (Melit.), M.Ent. (Melit.), Ph.D. (Melit.), M.I.E.E.E.

RESEARCH SUPPORT OFFICERS

Mr Russell Farrugia, B.Eng. (Hons)(Melit.), M.Sc.(Melit.) (Research Support Officer II)
Mr Barnaby Portelli, B.Eng. (Hons)(Melit.), M.Sc.(Melit.) (Research Support Officer II)
Mr Matthew Meli, B.Sc. (Hons)(Melit.), M.Sc. (Melit.) (Research Support Officer II)

ADMINISTRATIVE & TECHNICAL STAFF

Ms Alice Camilleri (Administrator I)
Inġ. Francarl Galea, B.Eng. (Hons.),M.Sc.(Eng.) (Senior Systems Engineer)

Research Areas:

- Analogue and Mixed Mode ASIC Design
- Radio Frequency Integrated Circuits
- Embedded Systems
- Biotechnology Chips
- Micro-Electro-Mechanical Systems (MEMS)
- Quantum Nanostructures
- System-in-Package (SiP)
- System-on-Chip (SoC)
- Accelerator Technology
- Microfluidics

DEPARTMENT OF ARTIFICIAL INTELLIGENCE

PROFESSOR

Professor Matthew Montebello, B.Ed. (Hons)(Melit.), M.Sc. (Melit.), M.A. (Ulster), Ph.D. (Cardiff), Ed.D. (Sheff.), SMIEEE (Head of Department)

ASSOCIATE PROFESSORS

Professor Alexiei Dingli, B.Sc.I.T. (Hons.) (Melit.), Ph.D. (Sheffield), M.B.A (Grenoble)

SENIOR LECTURERS

Dr Joel Azzopardi, B.Sc. (Hons.) (Melit.), Ph.D. (Melit.)

AFFILIATE SENIOR LECTURER

Mr Michael Rosner, M.A. (Oxon.), Dip.Comp.Sci.(Cantab.)

LECTURERS

Dr Charlie Abela, B.Sc. I.T. (Hons)(Melit.), M.Sc. (Comp.Sci.)(Melit.),Ph.D.(Melit.)
Dr Claudia Borg ,B.Sc. I.T. (Hons.) (Melit), M.Sc. (Melit.), Ph.D. (Melit.)
Dr Vanessa Camilleri, B.Ed. (Hons.)(Melit.), M.IT (Melit.), Ph.D. (Cov)
Dr Josef Bajada, B.Sc. I.T. (Hons)(Melit.), M.Sc. (Melit.), M.B.A.(Henley), Ph.D. (King`s)
Dr Ingrid Vella, B.Eng. (Hons)(Melit.), M.Sc. (Imperial), D.I.C., Ph.D. (Nott.), M.B.A. (Lond.)

ASSISTANT LECTURERS

Mr Kristian Guillaumier, B.Sc. I.T. (Hons.) (Melit.), M.Sc. (Melit.)
Mr Dylan Seychell, B.Sc. I.T. (Hons.) (Melit.), M.Sc. (Melit.), GSMIEEE

RESEARCH SUPPORT OFFICERS

Mr Mark Bugeja, B.Sc. (Hons.) Creative Computing (Lond.), M.Sc. AI (Melit.) (Research Support Officer II)
Mr Luca Bondin, B.Sc. IT (Hons) (Melit.), M.Sc. AI (Melit.) (Research Support Officer II)
Mr Foaad Haddod, B.Sc. (Al-Jabal Al Gharbi), M.Sc. AI (Melit.) (Research Support II)

ADMINISTRATIVE & TECHNICAL STAFF

Ms Francelle Scicluna, (Administration Specialist) B. W.H.R (Hons.) (Melit.)

Actual research being done

- Title: Maltese Speech Recognition (MASRI) and Maltese Speech Synthesis
  Area: Speech Processing
- Title: EnetCollect – Crowdsourcing for Language Learning
  Area: AI, Language Learning
- Title: Language in the Human-Machine Era
  Area: Natural Language Processing
- Title: Language Technology for Intelligent Document Archive Management
  Area: Linked and open data
- Title: Learning Analytics, Ambient Intelligent Classrooms, Learner Profiling
  Area: ICT in Education
- Title: Medical image analysis and Brain-inspired computer vision
  Area: Intelligent Image Processing
- Title: MyOcean Follow-On, MEDESS4MS, and Calypso 2 projects
  Area: Down-stream services
- Title: GBL4ESL
  Task: Creation of digital resources for educators using a Game Based Learning Toolkit
- Title: eCrisis
  Task: Creation of framework and resources for inclusive education through playful and game-based learning

- Title: Smart animal breeding with advanced machine learning techniques
  Area: Predictive analysis, automatic determination of important features
- Title: Real-time face analysis in the wild
  Area: Computer vision
- Title: RIVAL; Research in Vision and Language Group
  Area: Computer Vision/NLP
- Title: Maltese Language Resource Server (MLRS)
  Area: Natural Language Processing
- Task: Research and creation of language processing tools for Maltese
- Title: Walking in Small Shoes: Living Autism
  Area: Virtual Reality
  Task: Recreating a first-hand immersive experience in autism
- Title: Augmenting Art
  Area: Augmented Reality
  Task: Creating AR for meaningful artistic representation
- Title: Morpheus
  Area: Virtual Reality
  Task: Personalising a VR game experience for young cancer patients
- Title: Notarypedia
  Area: Knowledge Graphs and Linked Open Data
- Title: Smart Manufacturing
  Area: Big Data Technologies and Machine Learning
- Title: Analytics of patient flow in a healthcare ecosystem
  Area: Blockchain and Machine Learning

An updated list of concrete areas in which we have expertise to share/offer

- AI, Machine Learning, Adaptive Hypertext and Personalisation
- Pattern Recognition and Image Processing
- Web Science, Big Data, Information Retrieval & Extraction, IoT
- Enterprise Knowledge Graphs
- Agent Technology and Ambient Intelligence
- Drone Intelligence
- Natural Language Processing/Human Language Technology
- Document Clustering and Scientific Data Handling and Analysis
- Intelligent Interfaces, Mobile Technologies and Game AI
- Optimization Algorithms
- AI Planning and Scheduling
- Constraint Reasoning
- Reinforcement Learning
- AI in Medical Imaging Applications (MRI, MEG, EEG)
- Gait Analysis
- Machine Learning in Physics
- Mixed Realities

Other areas which we are interested in can be found at the below link:

https://www.um.edu.mt/ict/ai/ourresearch/departmentprojects

DEPARTMENT OF COMPUTER INFORMATION SYSTEMS

ASSOCIATE PROFESSOR
Professor Ernest Cachia, M.Sc.(Kiev), Ph.D.(Sheff.) (Head of Department)

SENIOR LECTURERS
Dr John Abela, B.Sc.(Hons.), M.Sc., Ph.D.(New Brunswick), I.E.E.E., A.C.M.
Dr Lalit Garg, B.Eng.(Barkt), PG Dip. I.T.(IIITM), Ph.D.(Ulster)
Dr Colin Layfield, B.Sc. (Calgary), M.Sc.(Calgary), Ph.D.(Leeds)
Dr Peter A. Xuereb, B.Sc.(Eng.)(Hons.)(Imp.Lond.), A.C.G.I., M.Phil.(Cantab.), Ph.D.(Cantab.)
Dr Joseph Vella, B.Sc., Ph.D.(Sheffield)

VISITING SENIOR LECTURERS
Dr Vitezslav Nezval, M.Sc.(V.U.T.Brno),Ph.D.(V.A.Brno)
Mr Rodney Naudi, B.Sc., M.Sc.(Eng.)(Sheff.)

LECTURERS
Dr Conrad Attard, B.Sc.(Bus.&Comp.), M.Sc., Ph.D.(Sheffield) (Deputy Dean of Faculty)
Dr Michel Camilleri, B.Sc., M.Sc., Dip.Math.&Comp., Ph.D (Melit.)
Dr Clyde Meli, B.Sc., M.Phil, Ph.D (Melit)
Dr Christopher Porter, B.Sc.(Bus.&Comp.), M.Sc. , Ph.D.(UCL)

VISITING ASSISTANT LECTURERS
Inġ. Saviour Baldacchino, B.Elec.Eng.(Hons.), M.Ent., D.Mgt.
Mr Norman Cutajar, M.Sc. Systems Engineering

ASSISTANT LECTURER
Mr Joseph Bonello, B.Sc.(Hons)IT(Melit.), M.ICT(Melit.)

ASSOCIATE ACADEMIC
Mr Anthony Spiteri Staines, B.Sc., M.Sc., A.I.M.I.S., M.B.C.S.

ADMINISTRATIVE STAFF
Ms Shirley Borg, (Administration Specialist)
Ms Lilian Ali, (Administrator I)

Research Areas:

Software Engineering

- Computational complexity and optimisation
- Integrated risk reduction of information-based infrastructure systems
- Model extraction (informal descriptions to formal representations)
- Automation of formal programming syntax generation
- Automation of project process estimation
- High-level description language design
- Distributed computing systems and architectures
- Requirements engineering - methods, management and automation
- System development including real-time scheduling, stochastic modelling, and Petri-nets
- Software testing, information anxiety and ergonomics

Data Science and Database Technology

- Data integration and consolidation for data warehousing and cloud services
- Database technology, data sharing issues and scalability performance
- Processing of streaming data
- Data analysis and pre-processing
- Predictive modelling
- Data warehousing and data mining: design, integration, and performance
- Big data and analytics

- Search and optimization
- Business intelligence
- Data modelling including spatial-temporal modelling
- Distributed database systems
- Missing data analysis
- Information retrieval

Human-Computer Interaction

- Human-Computer Interaction (HCI)
- Understanding the User Experience (UX) through physiological and cognitive metrics
- Human-to-instrumentation interaction in the aviation industry
- User modelling in software engineering processes
- Human-factors and ergonomics
- Accessibility, universal design and accessible user agents
- Advancing assistive technologies (multi-modal interaction)
- Affordances and learned behaviour
- The lived experience of information consumers
- Information architecture

Bioinformatics, Biomedical Computing and Digital Health

- Gene regulation ensemble effort for the knowledge commons
- Automation of gene curation; gene ontology adaptation
- Classification and effective application of curation tools
- Pervasive electronic monitoring in healthcare
- Health and social care modelling
- Missing data in healthcare records
- Neuroimaging
- Metabolomics
- Technology for an ageing population
- Education, technology and cognitive disabilities (e.g. augmented reality)
- Assistive technologies in the context of the elderly and individuals with sensory and motor impairments in institutional environments
- Quality of life, independence and security - investigating the use of robotic vehicles, spoken dialogue systems, indoor positioning systems, smart wearables, mobile technology, data-driven systems, machine learning algorithms, optimisation and spatial analytic techniques

Applied Machine Learning, Computational Mathematics and Statistics

- Applicative genetic algorithms and genetic programming
- Latent semantic analysis and natural language processing
- Heuristics and metaheuristics
- Stochastic modelling & simulation
- Semantic keyword-based search on structured data sources
- Application of AI and machine learning to business and industry
- Application of AI techniques for operational research, forecasting and the science of management
- Application of AI techniques to detect anomalies in the European Electricity Grid
- Knowledge discovery
- Image Processing (deconvolution)
- Image super-resolution using deep learning techniques
- Optimization of manufacturing production lines using AI techniques
- Square Kilometre Array (SKA) Tile Processing Module development
- Spam detection using linear genetic programming and evolutionary computation
- Scheduling/combinatorial optimisation
- Traffic analysis and sustainable transportation
- Automotive cyber-security

Fintech and DLT

- Automatic Stock Trading
- Distributed Ledger Technologies

FACULTY OFFICE

Ms Nathalie Cauchi, Dip.Youth&Comm.Stud.(Melit.), H.Dip.(Admin.&Mangt.) (Melit.), M.B.A.(A.R.U.,UK) (Manager II)
Ms Michelle Agius, H.Dip.(Admin.&Mangt.)(Melit.) (Administrator I)
Mr Rene' Barun, BA (Hons.) Philosophy (Melit), (Administrator I)
Mr Anthony Buhagiar, (Senior Administrator)
Ms Therese Caruana, (Administrator I)
Ms Anabel Decesare, (Administration Specialist)

SUPPORT STAFF

Mr Patrick Catania A.I.M.I.S. (Senior IT Officer I)
Mr Paul Bartolo (Senior Beadle)
Ms Melanie Gatt (Beadle)
Mr Raymond Vella (Technical Officer II)

YOU'RE PAYING TOO MUCH FOR YOUR CLOUD SERVICES

Let's talk about cost optimization

CLOUD@ICTSOLUTIONS.COM.MT

*Over 90% of Azure users spend too much on their public cloud. But we can help.

ICT SOLUTIONS

Microsoft Partner
Microsoft
2020 Partner of the Year Winner
Malta

Microsoft Partner
Microsoft
2019 Partner of the Year Winner
Malta

DISCOVER THE WORK TAKING PLACE AT THE

Faculty of Information & Communications Technology

Since its inception, the Faculty of ICT at the University of Malta has worked on creating an environment where students can achieve their full potential and where industry can find a trusted partner to collaborate with.

Human beings are an ingenious species: we have, after all, found solutions to so many problems that the lives of those inhabiting the earth today look almost nothing alike those of our ancestors who lived a mere hundred years ago.

This development has been crucial for a number of things we take for granted including longer life expectancy, better communication, more comfortable living, and faster travel. Yet, while many of these past successes were down to analogue development, the future requires us to go digital.

Today, Information and Communication Technology (ICT) is at the forefront of the drive to give humans better ways to communicate, diagnose illness and even extend abilities. And behind each new breakthrough is research conducted by universities, industries, or a consortium made up of the two.

Here at the Faculty of ICT within the University of Malta (UM), we are among the most prolific researchers in the areas of ICT on the Maltese Islands, with our lecturers' and students' research covering a wide spectrum of disciplines. More importantly, on top of collaborating with one another, the five departments that make up the Faculty also work with numerous local and foreign business entities, research centres and universities.

This ensures that our students work in an environment in which they

are exposed to the latest ideas and the most cutting-edge technology on the subject at hand. In fact, as an educational institution, this is part of our responsibility to ensure that we not only have professionals who are at the top of their game once they graduate, but that we also have citizens who are not limited in their way of thinking, in the way they look at problems, and in the way they seek to achieve solutions.

Over the past few years, the Faculty has indeed done its best to exceed the expectations of its students and the general public, which it ultimately serves. Students and lecturers of the Faculty of ICT have been given world-class opportunities including working closely with the Paul Scherrer Institute in Switzerland and the National Institute for Research in Digital Science and Technology in France, as well as training from global industry giants like Philips, Microsoft, ST Microelectronics, and Google.

Industry, however, doesn't always just mean global tech giants. Through a collaboration with St Vincent de Paule Residence, for example, members of the faculty are currently working on a study on how pervasive electronic monitoring through the use of mobile applications could help caregivers and family members monitor people with dementia and increase their quality of life.

But that is still just the tip of the iceberg. Through our perseverance to form new collaborations with partners across Malta, Europe and beyond – as well as our students' and lecturers' determination – the Faculty now runs research on a huge number of topics.

In this publication, we look at more than ten different research projects currently taking place, which include teaching computers how to properly read images and the sentiment behind sentences, aiding people living with disabilities in their everyday lives, improving the way archaeologists conduct their research, and helping healthcare providers give a better and more seamless service.

Yet the publication only shows a small array of what is actually happening. Indeed, over and above the research you'll read about in this publication, members of the Faculty are also currently working on some other ground-breaking projects.

There is one study, for example, that is looking to improve the spatial resolution of images of Malta and Gozo's coastal areas so that we can predict marine currents more accurately, as well as aid divers and search and rescue operations in their work. Another study, meanwhile, is looking to create more secure digital communication in the Quantum Era – a project that is incredibly useful at a time when one of the biggest hindrances in development is our very real worry that our data might end up in the wrong hands.

> "The beauty of ICT is that all areas and industries need it"

Together, these projects show that the beauty of ICT is that all areas and industries need it. Whenever we are looking to innovate or to create something new, ICT is now an essential tool that enables and empowers change. And, while industry has a big say in how things develop, students can ultimately spearhead where we go.

But the work with business is still crucial and, thankfully, the Faculty of ICT has a great relationship with many local industries. This is something that is benefitting both parties. After all, industry in Malta may not always have the right structures to conduct research all by itself. Yet, at UM, we have these structures in place, and we also have the knowledge, the time and the willingness to commit to research, which can sometimes take years. Moreover, we continually invest in the infrastructure needed to conduct such research, which is an essential element in ensuring that we give all the research taking place the right environment to succeed.

All this is part of the reason why the Faculty continues to remain relevant to modern-day needs, which is why so many industries look to collaborate with us. But, as aforementioned, it's a two-way street and this collaboration aids our team, too.

Over the past few years, collaboration with local businesses has seen them share invaluable accumulated data sets (always within GDPR guidelines) with the Faculty, as well as best practices, technology and machinery.

But, more importantly, industry is seeing the benefit of working closely together with us, which is why it also provides training for our researchers, offers licensing tools, and even helps fund certain project.

Funding, of course, is extremely important when it comes to research. And it's for this reason that we are always working on finding new sources to ensure the Faculty has whatever it needs to continue thriving. There are many entities that we need to thank for that, including the Research, Innovation and Development Trust (RIDT), Interreg Europe, FinTech, numerous European Union funds, MCST and, of course, many industries and businesses. Indeed, it's through them that the Faculty's many research projects continue to happen, and we are extremely grateful for their support.

With this in mind, we feel the Faculty is adequately prepared to take on a whole new group of students, who will be joining those already reading for their Bachelor's and Master's degrees, as well as their PhDs. But, more than that, the Faculty believes that there has never been a better time for prospective students to follow their dreams by joining those already making waves in the world of ICT.

> The selection of degrees offered by the Faculty of ICT is indeed the biggest it's ever been.

The selection of degrees offered by the Faculty of ICT is indeed the biggest it's ever been. There are also a lot of resources students can make use of at this point in time: we have scholarships, funding for research, good teams in place, industries that support us, and even government help. There's no other way of putting it: it's a truly exciting time to join the Faculty and to explore the topic of your choice.

Those looking to start an undergraduate degree can choose between Artificial Intelligence, Communications & Computer Engineering, Computer Information Systems, Computer Science, and Microelectronics & Nanoelectronics. Meanwhile, those looking to read for a post-graduate degree have even more choice, with Master's in Telecommunications, Computer Science, Computer Information Systems, Microelectronics & Microsystems, Signal Processing, and Machine Learning & Human Language Science being among the many specialisations available. And if it's AI you're looking to read a Master's in, then we've even got scholarships

sponsored by the Malta Digital Innovation Authority. And, of course, there are even more opportunities at doctoral level.

As aforementioned, all these degrees will form the basis of what students can focus on, and each degree can be used to progress in various industries including healthcare, chip design, gaming, telecommunications, FinTech, aviation and archaeology, to mention but a few.

From our end, we promise to give you all the space you need to explore your ideas and to help you find a sector that interests you. Indeed, we invite prospective students to speak to us about their goals and visions so we can work together on making them a reality.

With all this in mind, we invite you to go through the rest of the publication to discover how broad the world of ICT truly is.

Discover the research projects mentioned above by following these links:

Research work that uses ICT solutions to help monitor people with dementia
https://researchtrustmalta.eu/research-projects/pem/

Research work on monitoring elderly at home
https://natiflife-project.eu/

Research work that uses satellite imagery and improves its resolution to monitor coastal areas
https://www.um.edu.mt/projects/sat-fire/

Research on secure communications
https://re-search.info/

Blockchain-based traceability of the wine supply chain

Gabriel Camilleri | **Supervisor: Prof. Gordon J. Pace**

Co-Supervisor: Dr Joshua Ellul

Course: B.Sc. (Hons.) Computing Science

Supply-chain systems, which handle the flow and transformation of goods, have been present since the inception of commercialisation. Over time, the management of these chains has been an increasing focus for companies in order to obtain the information necessary for maintaining and improving the performance of workers and the efficiency of production. Moreover, through ongoing analysis of, and tweaks to, the supply chain, one of the goals of supply-chain management is for the system to be modified to have more transparent, flexible, and robust properties.

This project focuses on the managerial process in the sector of wine distribution. Apart from the issues also present in regular supply chains, this sector also encounters issues specific to wine distribution, such as wine fraud (a). Thus, the salient characteristics of the issues being addressed through this study concerned the concentration of authority and the traceability of the commodities. The main objectives of this project were to explore the design options of the management of this supply chain in a smart-contract view, and to identify properties that could be formally verified.

The design of how the wine-distribution supply chain operates was acquired through a study of papers/articles related to the sector [2,3]. The entities and their relationship with other entities were defined. The main phases of the wine supply chain are as follows: the vineyard, where the grapes are produced; the manufacturer, who handles the winemaking process; and finally, the market, which is the end-buyer or consumer who relies on the distributor (who, in turn, is the link between consumer and winemaker). The flow and transformation of the items are considered through all the elements involved between the vignerons and the consumers.

Blockchain technology makes use of a digital distributed ledger of transactions accessible by shared nodes on a network in a peer-to-peer and decentralised manner. The transactions are immutable, secure, and reliable. The implementation of the wine supply chain system is built through the use of smart contracts. These are self-executing programs that have been fed with a set of requirements to run functionalities that could include transferring monetary values from one address to another [4]. The smart contracts in this project fall under the Ethereum cryptocurrency platform, and programmed using Solidity. The above-mentioned technologies have been utilized with a view to mitigate the pitfalls of centralised systems, reduce points in the system containing dependencies and safety risks as well as ensuring a system with increased transparency and better traceability.

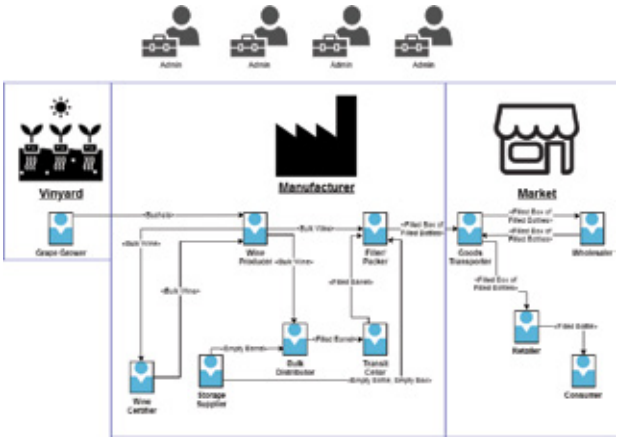

Figure 1. Overview of the wine supply chain

The project also included a run-time verification element. This is a type of formal verification, which entails an evaluation process to check specifications on the extent to which they satisfy well-formed requirements, and expected properties are upheld. Run-time verification tests properties while the system is still executing. Thus, the properties are defined by the expected behaviour of the smart contract, proving that the resulting behaviour indicates that the smart contract satisfies the intended goals.

Finally, the tool of ContractLarva verification tool was used to generate an updated and safe smart contract by providing the initial smart contract and the properties of the smart contract, which are expressed in dynamic event automata [1]. By comparing the cost to execute the initial smart contract with the cost for running the modified version, it would be determined which run-time verification tests are important enough to still be included in the system.

a. - Wine Fraud is the falsification of product details in wine commerce.

References/Bibliography:

[1] K. Biswas, V. Muthukkumarasamy, and W. L. Tan, “Blockchain based wine supply chain traceability system,” 2017.

[2] M. G. Cimino and F. Marcelloni, Enabling traceability in the wine supply chain. Springer, 2012, pp. 397–412.

[3] P. Tsao, “Blockchain 2.0 and Ethereum [Blockchain Basics Part 3],” 2018. [Online]. Available: https://medium.com/xpa-2-0/blockchain-2-0-and-ethereum-blockchain-basics-part-3- 362eb3561b4e

[4] S. Azzopardi, J. Ellul, and G. J. Pace, “Monitoring smart contracts: Contractlarva and open challenges beyond,” 2018, pp. 113–137.

Credit card fraud detection with oversampling

Julian Demicoli | **Supervisor: Dr John Abela**

Course: B.Sc. IT (Hons.) Computing and Business

Card-based payments are increasingly becoming the standard payment method by consumers. Indeed, between 2017 and 2018 global expenditure attributed to card payments grew by 17.7% to $40.582 trillion.

Many industries rely heavily on card-based payments as an efficient means of collecting money from consumers, as is the case with e-commerce. As the number of annual transactions made through the various types of payment cards increases, losses due to fraud are also on the rise and these are expected to amount to $35.67 billion globally by 2023 [1]. Given the quantity of transactions being processed on a daily basis by financial institutions, and the substantial losses being incurred due to fraud, the said institutions must implement fraud-detection systems that are cost-effective, automated, and offering high accuracy with minimal human intervention. This issue has not gone unnoticed by researchers and strides have been made in early-fraud-detection systems, which make use of machine learning techniques. However, the imbalanced distribution between fraudulent (minority) and non-fraudulent (majority) is challenging for many traditional learning algorithms, which are ill-suited to handle such large-class imbalances [2]. So much so, that models trained on these datasets tend to be biased towards the majority class and yet still achieve high accuracy scores because the minority class is ultimately a very small percentage of the dataset.

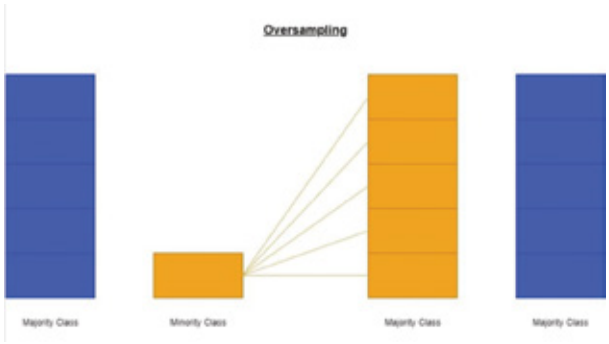

Figure 1. Oversampling

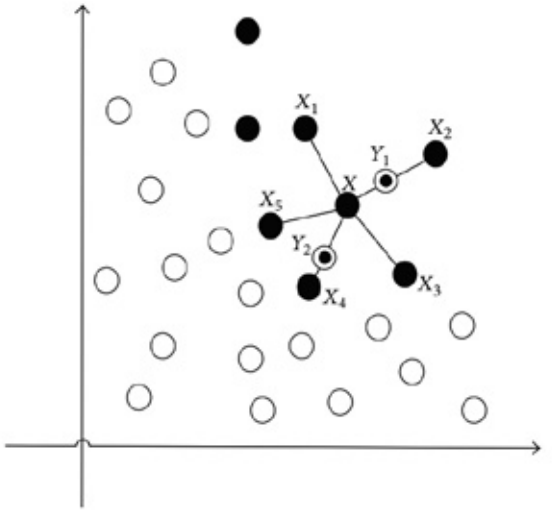

Figure 2. Synthetic Minority Over-Sampling Technique (SMOTE) [4]

In this final year project, a data-level approach has been used to overcome class imbalance by incorporating oversampling techniques that make use of synthetic data. Within this domain, a technique known as Synthetic Minority Over-Sampling Technique (SMOTE) [3] has been used extensively, yet the development of SMOTE has spawned may variants of this technique, which have not been examined as extensively. This project evaluates other popular variants of the original technique, in conjunction with machine learning techniques that have been considered in relevant literature to perform well in this domain. These techniques are: XGBoost, random forest and Gaussian Naïve Bayes. Finally, the project sought to highlight the most effective algorithm for card fraud detection purposes.

References/Bibliography:

[1] "Card Fraud Losses Reach $27.85 Billion," Nilson Report – Card Fraud Losses Reach $27.85 Billion. [Online]. Available: https://nilsonreport.com/mention/407/1link/. [Accessed: 03-Aug-2020].

[2] G. M. Weiss and F. Provost, "Learning When Training Data are Costly: The Effect of Class Distribution on Tree Induction," Journal of Artificial Intelligence Research, vol. 19, pp. 315–354, 2003.

[3] N. V. Chawla et al, "SMOTE: synthetic minority over-sampling technique," Journal of Artificial Intelligence Research, vol. 16, pp. 321-357, 2002.

[4] X. Yonghua, L. Yurong, and F. Qingqiu, "Research Article Imbalanced Data Sets Classification Based on SVM," hindawi.com. [Online]. Available: https://www.hindawi.com/ journals/ddns/2015/562724/. [Accessed: 05-Aug-2020].

Using Business Intelligence and machine learning on ERP data to improve business performance

Jesmar Scicluna | **Supervisor: Dr Peter Albert Xuereb**

Co-Supervisor: Mr Anthony Spiteri Staines

Course: B.Sc. IT (Hons.) Computing and Business

Companies today have grown rich in data [1]. Moreover, by analysing their data, businesses could capture and create value from it [2]. Therefore, since data analysis would allow businesses to monitor and improve their performance by turning their data to information, business-data analysis is of paramount significance, all the more so since it would also enable businesses to create predictions through machine learning [2].

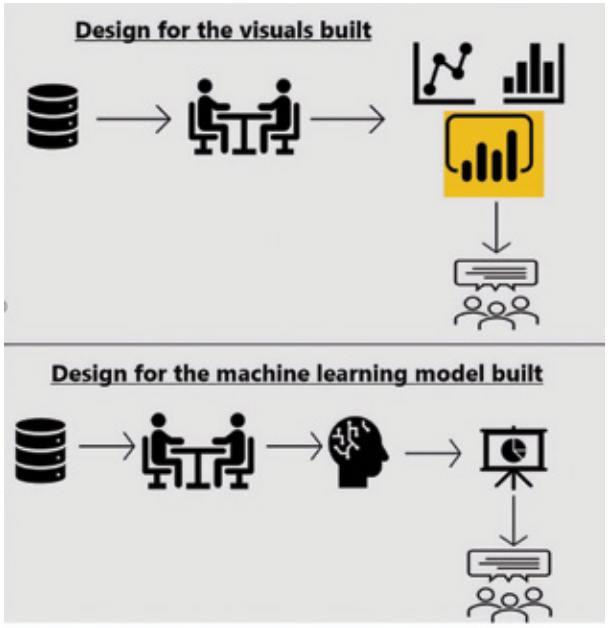

Figure 1. The design processes adapted in this study

The aim of this study was to demonstrate this by making use of business intelligence (BI) through data visualisation, and by creating predictions through a machine learning model; this is to enable companies to improve their business performance by making more informed decisions with the use of these technological tools. This study, therefore, assesses the usability and value that such technologies could deliver to companies seeking to make better use of their data.

To this end, the study benefitted from the participation of a local company, which offered its enterprise resource planning (ERP) data related to the company's sales. Various visuals/dashboards were built on stakeholder input within the company, and the data related to sales and sale opportunities. Moreover, a supervised multi-class classification machine learning model was built on the basis of the sale-opportunities dataset. This dataset was obtained from the company's customer relationship management (CRM) system to predict the class label for each sale opportunity (refer to Figure 1 for the design process adapted in this study).

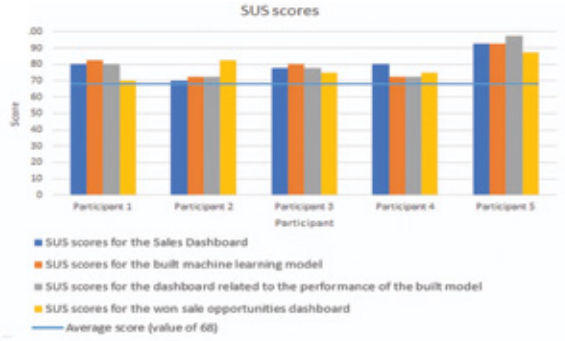

Figure 2. System usability scale (SUS) scores obtained by each participant involved in the evaluation process, where each score is higher than the general SUS average score (value of 68)

The visualisations and the machine learning model built were subsequently submitted for evaluation to identify the potential value of such technologies. Various employees, working at different levels within the company, were involved in the evaluation process for the purpose of gathering feedback and assessing usability of the artefacts developed. As a result of the evaluation process, these employees suggested also extending BI and machine learning to areas other than sales.

Given that positive feedback and good usability scores (Figure 2) were achieved, visualisation and prediction were classified as effective and important tools to enhance a company's business performance. These were acknowledged as useful tools in helping a company to improve its decision-making process by having timely access to the relevant information [3].

References/Bibliography:

[1] Anonymous, "Business Intelligence: Better Use of Data to Make Smart Business Decisions: CRM", Customer Relationship Management, 2017. Available: https://search-proquest- com.ejournals.um.edu.mt/docview/1973924342?accountid=27934.

[2] Anonymous, "The Value of Business Intelligence in ERP Systems: [Source: NoticiasFinancieras]", NoticiasFinancieras, 2015. Available: https://search-proquest- com.ejournals.um.edu.mt/docview/1722395561?accountid=27934.

[3] B. Hočevar and J. Jaklič, "Assessing Benefits of Business Intelligence Systems – A Case Study", Management: Journal of Contemporary Management Issues, vol. 15, no. 1, 2015. [Accessed 4 December 2019].

Meltdown and Spectre Vulnerabilities, and Performance Issues of Related Patches

Kevin Agius **Supervisor: Dr Clyde Meli**

Course: B.Sc. IT (Hons.) Software Development

Meltdown and Spectre are hardware vulnerabilities at the microarchitectural level of a central processing unit (CPU), allowing an attacker to manipulate various performance-enhancement techniques in order to extract sensitive data from cache memory [1,2]. The vulnerabilities make direct use of speculative execution, branch prediction and out-of-order execution, whereby the attacker injects malicious code in order to obtain data from the CPU cache memory [1,2]. Software patches for various operating systems and system kernels have been developed and were distributed in order to mitigate the severity of an attack [3]. Since the vulnerabilities in question are relatively new, a watertight preventive solution is yet to be identified. However, these patches have been deployed on Windows and a number of other operating systems have proved to be successful in slowing down the memory-reading process of the attack [4].

Gauging the performance impact of these patches entails the utilisation of benchmarking tools. Therefore, a number of benchmarking tools were tested in order to establish which was the most relevant for this study. The benchmarks were filtered out of the study according to their configuration capabilities and their relevance in terms of any changes that the patches implemented, mainly: out-of-order execution, branch prediction and speculative execution [1,2]. This criterion was adopted primarily because the changes that were introduced related directly to these functionalities.

It was concluded that the benchmarking tool 7-Bench5 would be the most suitable, since it makes direct use of the aforementioned operations, in particular branch prediction and out-of-order execution. The benchmarking tool was implemented on a wide variety of modern-generation processors and system frameworks, such as virtual machines and cloud-based systems, in order to analyse the impact that such patches have caused across various platforms. The benchmarking tool was executed on the systems in various scenarios, covering benchmarks with all the patches enabled and disabled, and for each patch individually.

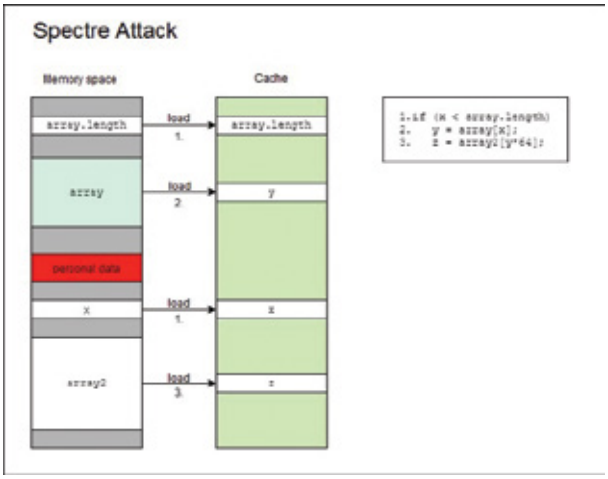

Figure 1. Example of Spectre attack

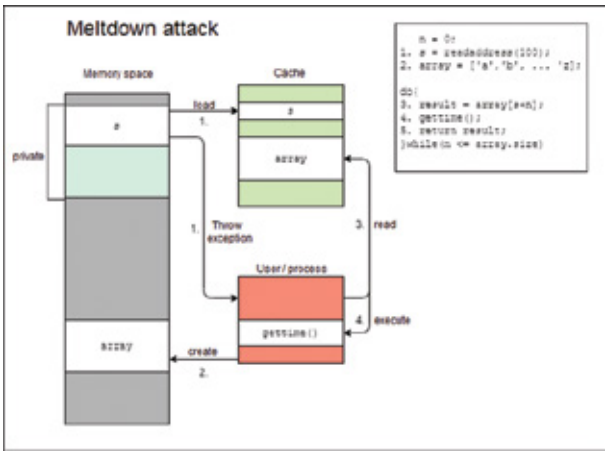

Figure 2. Example of Meltdown attack

The results obtained from the said benchmarks were then analysed in order to determine any resulting performance-degradation impacts by comparing the results with the default scenario, where all the patches available were enabled. Finally statistical analysis was used to test the hypothesis set for this study.

References/Bibliography:

[1] M. Lipp et al., 'Meltdown: Reading Kernel Memory from User Space', in 27th USENIX Security Symposium (USENIX Security 18), 2018

[2] P. Kocher et al., 'Spectre Attacks: Exploiting Speculative Execution', in 40th IEEE Symposium on Security and Privacy (S&P'19), 2019.

[3] https://www.csoonline.com/article/3293015/patching-windows-for-spectre-and-meltdown-a- complete-guide.html

[4] O. Alhubaiti and E. M. El-Alfy, "Impact of Spectre/Meltdown Kernel Patches on Crypto-Algorithms on Windows Platforms," 2019 International Conference on Innovation and Intelligence for Informatics, Computing, and Technologies (3ICT), Sakhier, Bahrain, 2019, pp. 1-6, doi: 10.1109/3ICT.2019.8910282.

[5] https://www.7-cpu.com/utils.html

Investigating MR interface patterns for just-in-time QA information

Matthew Cachia | **Supervisor: Dr Chris Porter**

Co-Supervisor: Dr Mark Micallef

Course: B.Sc. IT (Hons.) Software Development

Software testing is the predominant methodology used by the software development industry to build confidence in a software product prior to release. The ever-increasing complexity and volatility of requirements makes the task of software testers considerably difficult. This project builds upon previous work to investigate how mixed-reality design patterns, along with context-relevant just-in-time information, could be used to improve the performance of software testers.

Following an extensive review of design patterns related to mixed-reality systems, a subset of these was selected and incorporated into the design and implementation of a prototype utilising the Microsoft Hololens headset. To gain insight into the extent of the effect of this process on the performance of software testers, a qualitative study based on thematic analysis was undertaken. For the said study, experts from both fields (i.e. testing and mixed reality) participated in a series of review sessions using the prototype, showcasing how the approach could be used to enable testers to organise a virtual workspace in mixed reality, and how this workspace would then show just-in-time information in relation to whatever the tester was doing at any point in time. Subsequently, the participants were interviewed on how such an application could potentially improve and affect the workload of testers, as well as impact the industry as a whole.

Five key themes emerged from the interviews conducted with software-testing specialists, namely: 'Aiding the tester', 'Better understanding of data', 'Impact on the software industry', 'Pricing and integration', and 'Potential software features'. From the interviews with mixed-reality experts, the following three themes emerged: 'Interaction', 'Additional guidance', and 'Developers adapt'.

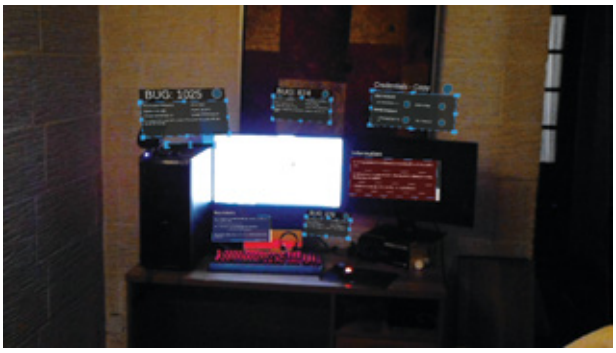

Figure 1. The application at work

The findings extracted from the interviews with the software testers highlighted how such a technology could not only improve a tester's workflow, but also have a positive impact on the industry at large. However, it was also noted that the current high pricing and integration costs hindered the progress of the technology in question. The findings resulting from feedback obtained from mixed-reality specialists underlined the importance of interaction between mixed-reality applications towards improving the experience of both new and advanced users.

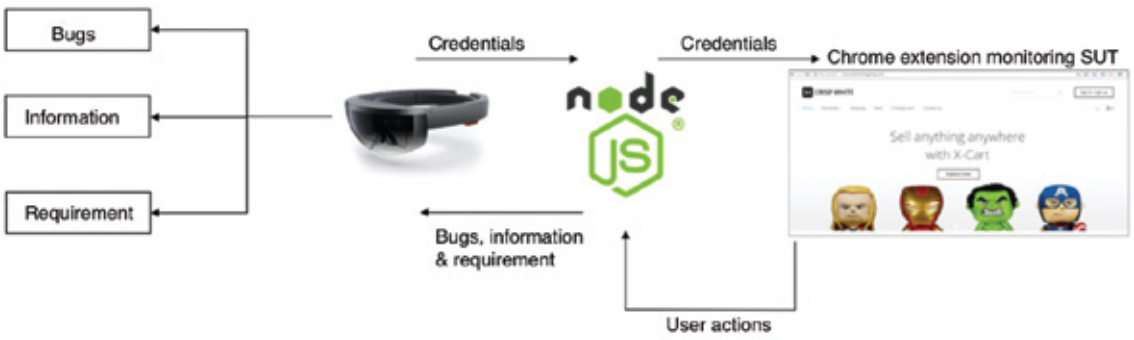

Figure 2. Overview of the system

Demand prediction for shared mobility services using time series modelling

Rudi Camilleri | **Supervisor: Mr Tony Spiteri Staines**

Co-Supervisor: Dr Peter Albert Xuereb

Course: B.Sc. IT (Hons.) Software Development

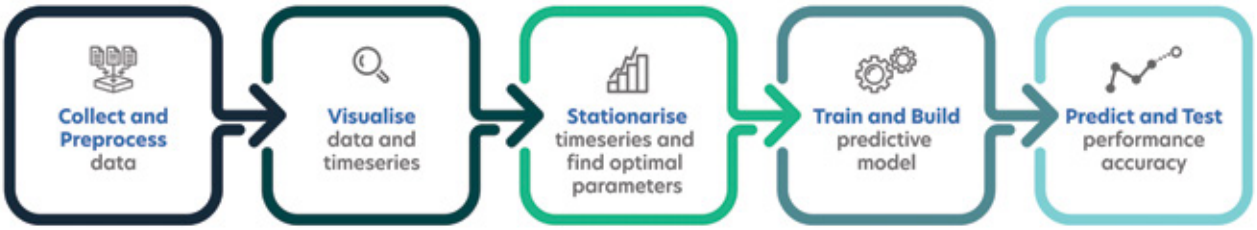

Figure 1. The methodology used in this study

People are becoming increasingly interested in new mobility solutions, such as shared mobility, which are both cheaper and environmentally friendly when compared to traditional transportation methods.

According to multiple reports, it is estimated that by 2025, the shared-mobility market worldwide will be worth over €550 billion [1]. However, as promising as ride-hailing mobility might be, it also has its challenges, with longer passenger waiting times, longer trips due to different pickup points and vacant vehicles driving around aimlessly waiting for a request [2]. The main objective of this project was to analyse and investigate the possibilities of optimising shared mobility using historical data and multiple predictive techniques. A further goal of the research was to be able to predict the total number of generated requests per hour for the near future. Applying the latest ICT technologies to the transportation industry contributes to the development of better transport infrastructure, with safer roads and better traffic management.

The accompanying diagram outlines the methodology used throughout the study, starting from data collection, pre-processing, and data visualisation. In the third step, where necessary, the time series is differentiated into a stationary time series as well as finding the optimal model-parameter combinations. This is followed by training and building the forecasting model. Finally, the trained model can predict future values which will be validated by comparing them with the testing dataset.

Using the data compiled by a Maltese ride-hailing company, this study first investigates where and how pickup requests were made. The data was subsequently examined for any time-series patterns [3]. The three main patterns considered to be the pillars for any series are: trend, seasonality and cyclic. When the 'requests' data is visualised geographically, it can be noted that most of these requests were booked around touristic areas, mainly Sliema and Valletta. From the results obtained, it could be implied that exogenous or independent data such as temperature and public holidays do not affect the predictive model. Such models have been validated by splitting the data between a training set and a testing set. The purpose of such sets is to fit the model on the training set and compare the outcome with the actual observed values of the latter set. As a result of the seasonality in the time series, the optimum models were the Holt- Winters model [4], the seasonal autoregressive integrated moving average (SARIMA) [5] and Facebook Prophet [6]. Using different performance metric calculations, the results conclude that amongst all models, the accuracy of the Holt-Winters model outperforms other models with an overall mean absolute error of 8.039 and root mean squared error of 11.159.

Finally, the study suggests that some models do not work as expected due to the lack of yearly seasonality. With this in mind, a report generator has been developed to forecast data on a real-time basis using the three above-mentioned predictive models.

References/Bibliography:

[1] "Mobility Scooter Market - Global Industry Analysis, Share, Forecast - 2024." [Online]. Available: https://www.transparencymarketresearch.com/mobility-scooter-market.html

[2] B. Schaller, "The new automobility: Lyft, Uber and the future of American cities," 2018.

[3] V. Shah, "A Comparative Study of Univariate Time-series Methods for Sales Forecasting," 2020.

[4] P. R. Winters, "Forecasting sales by exponentially weighted moving averages," Manage. Sci., vol. 6, no. 3, pp. 324–342, 1960.

[5] G. E. P. Box, G. M. Jenkins, and G. C. Reinsel, Time series analysis: forecasting and control, vol. 734. John Wiley & Sons, 2011

[6] S. J. Taylor and B. Letham, "Forecasting at scale," Am. Stat., vol. 72, no. 1, pp. 37–45, 2018.

Rethinking the bug-tracking paradigm

Kieran Cauchi | **Supervisor: Dr Chris Porter**

Co-Supervisor: Dr Mark Micallef

Course: B.Sc. IT (Hons.) Software Development

In modern development processes, the bug-tracking tool is one of the most important means of tracking any tasks, features or bugs active in a project. Such information is crucial, not only for maintaining developer workflow, but also to make informed decisions as to whether a product should go live in its state at a given point in time. When a bug is identified, the reporter creates a report using a bug-tracking tool. Bug reports typically include general information such as a summary of the issue, its expected due date, and a description. Once the report has been created, teams could rank and prioritise it as deemed fit, and assign it to individuals for resolution.

The main issue with the above approach in tracking project progress is that the quality of bug reports could vary depending on the person reporting the bug, their knowledge of the system, their cognitive workload, and other contextual factors. This project sought to study prevalent issues with bug-tracking systems, and to propose an improvement on existing systems.

The research was carried out in two phases. Firstly, a literature review was undertaken in order to gain insight into prevalent issues with bug tracking. The literature review was complemented with a survey amongst practitioners to identify the main complaints concerning bug-tracking systems. Results showed that, according to the participants, the main critical issues were the lack of important information in bug reports, and miscommunication between bug reporters and developers.

Taking the above into account, a Jira plug-in was designed and implemented. This plug-in departed from the 'compulsory field' approach generally adopted in improving bug quality, and instead analysed bug description content in real time to provide recommendations for improving upon what was being noted. As a visual aid, a suitability meter (akin to password-strength meters) was also developed. This provided an overall rating indicator regarding the quality of the bug report being created.

The plug-in was trialled by practitioners, and the resulting bug reports were reviewed and rated by a panel of two experts. The latter rated a mixture of bug reports, unaware which were the 50% written without the aid of the plug-in, and which were the 50% written with the aid of the plug-in. The results obtained from the expert reviews suggest that the plugin helped generate better bug reports, as rated in terms of completeness, usefulness and overall quality.

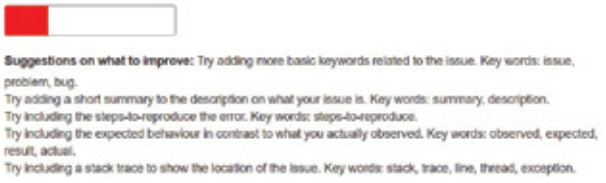

Figure 2. The report-quality meter and suggestions on what to improve in the report

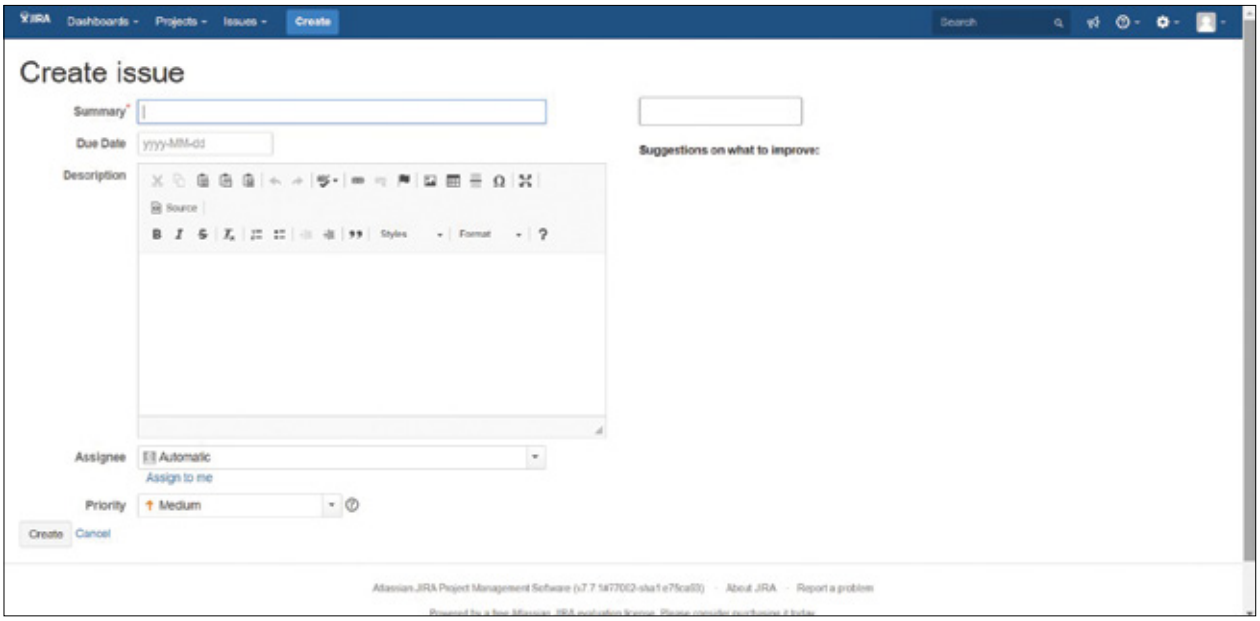

Figure 1. The 'Create issue' page

Investigating the use of genetic algorithms in automated test-case generation

Cristina Gatt **Supervisor: Dr Mark Micallef**

Course: B.Sc. (Hons.) Computing Science

Search-based Software Testing (SBST) is the reformulation of the test case generation problem as a search problem encountered when seeking to find the optimal feasible solution using meta-heuristic optimization techniques such as genetic algorithms [1]. Genetic Algorithms (GAs), inspired by the theory of natural selection. Indeed, they reflect the process of nature where the fittest individuals are more likely to reproduce and create offspring while allowing less fit individuals to reproduce in order not to kill diversity.

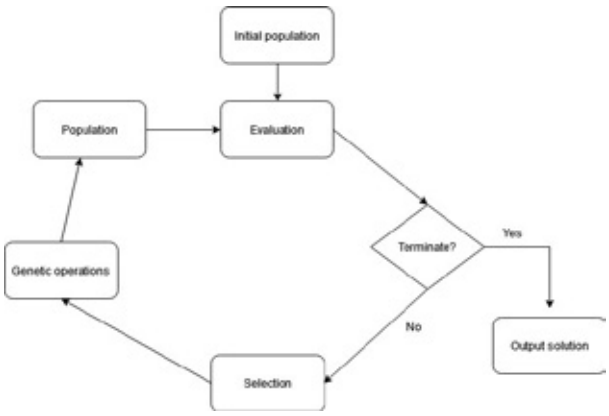

Figure 1. Flowchart outlining the genetic algorithm

GAs manipulate a population of candidate solutions to an optimization problem. Selection provides the driving mechanism for the best solutions to survive. Each solution is assigned a fitness value that reflects its effectiveness when compared with other solutions related to the same problem. The higher its' fitness, the higher its' chances to survive and reproduce. Crossover and Mutation techniques are used to simulate the recombination of genetic material [2].

Due to a paradigm shift towards more decoupled microservices -based architectures, the access of features of an application from another application have become common practice using REST APIs . This study, set out to investigate the use of GAs in automatic test-case generation for a non-trivial REST API library, focusing on the effect of the fitness function and the state of the database. To investigate the effect of the fitness function, two fitness functions were developed. The first was based on code coverage. In a bid to improve the obtained results, the concept of journey factors and tours as used in exploratory testing were introduced in a second fitness function. To investigate the effect of the state of the database, this was reset a) after every iteration and b) after every test case in different experiments.

The experiments were carried out using an open-source REST API. In general, the GA did not surpass the results generated randomly in terms of code coverage. On evaluating this result, it was noted that the chosen REST API did not have a lot of interdependent calls. Nevertheless, the GA whose fitness function was influenced by an appropriate number of journeys, yielded better results than the GA with a fitness function based only on coverage. Furthermore, when evaluating the realism of the generated test cases, those generated by the GA compared better with the tests written by a human tester than those generated by the random approach.

The GA was affected by the state of the database, with the experiments in which the database was reset after every iteration converging much faster than the others. This lead to the conclusion that to achieve optimal results, the database should not be reset after each test case. In order that a GA could produce a good test suite, further work should be carried out such as experimenting with different REST APIs and finding the best possible parameters- possibly also refining the fitness functions.

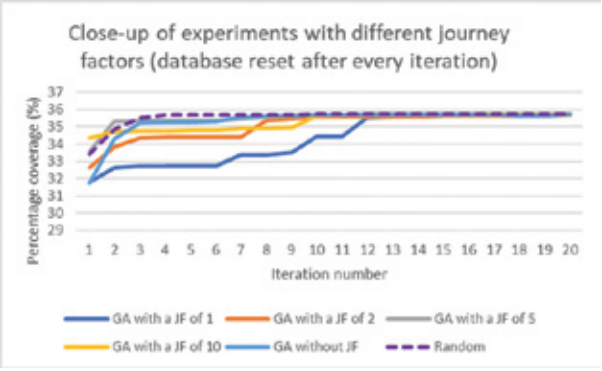

Figure 2. Results of experiments with the database being reset after every iteration

References/Bibliography:

[1] P. McMinn. Search-based software test data generation: a survey. Software testing, Verification and reliability, 14(2):105-156, 2004.

[2] M. Srinivas and L. M. Patnaik. Genetic algorithms: A survey. computer, 27(6):17-26, 1994.

A framework for instrument panel layout construction

Chakotay Incorvaia **Supervisor: Prof. Ernest Cachia**

Course: B.Sc. IT (Hons.) Software Development

The goal of this study is to produce techniques to generate alternative user interface (UI) designs and layouts for systems with their components. This may refer to both existing UI layouts and those that are yet to be designed. This study will show that UI layouts may benefit from algorithmic reasoning. Consequently, there may be more than one UX/UI design that could fulfil the same role, whilst also increasing human-to-system interaction efficiency and additional optimal designs.

This study seeks to add to the possibilities for use of UX (user experience) and UI design and, among other benefits, reducing the time required for developing these changes – while continuing to comply with the relevant standards. [1]. It focuses specifically on hardware-oriented UIs, such as instrument panels found in manufacturing plants, transportation systems and control systems. However, the principle could also be applied to software interfaces such as tablets. This study has faced multiple challenges, given the difference between physical and virtual interfaces and the wide array of environments in which they are deployed. Therefore, the ideal focus for this study would be an instrument layout that represents the majority of properties of interfaces found in most domains, whilst also containing attributes that would be present in most systems [2]. Aircraft cockpits have been found to contain the widest range of instruments with different interfaces and functions – hence being a valid test subject for this study [3]. Another challenge was on how and by what means would a layout be re-constructed. Different layouts may be specifically intended to control certain processes and there are different algorithms that could be used to carry out such control [4]. This is dependent on the nature of the problem from the point of view of the various stakeholders with respect to various factors, such as efficient space utilisation, cost effectiveness, ease of use, among others. The algorithm that would prioritise certain characteristics required thorough research and subsequent reasoning. Another challenge encountered was that many UIs, such as those found in aircraft and power stations, are multi-user interfaces. This issue falls outside the scope of this study. Nevertheless, the underlying proposed solution resulting from this study is a system that would take an interface as an input, and produce multiple new designs using the same components of the provided design. The solution would be produced from algorithms that could break down the interface into smaller components, which would be re-arranged as necessary, in line with specifications. Drawing from model-based design, the model could be broken down by components, to which valued attributes could be assigned individually [5]. Thus, an algorithm could utilise the newly valued components to create a new design.

This study produced three separate programs, each of which highlights a different aspect of the system. One of these focuses on displaying attributes and data of a design, and another focuses on producing new designs to be displayed in an instrument cluster. The remaining program produces new designs to be displayed and oriented in a limited area. The data has been processed through a user-submitted data file built against a template. The result of this study suggests that this proof-of-concept is valid and that a UI could be modified and improved using this methodology in order to improve the design and, consequently, the development process. This study could be further improved with more research on additional algorithms and further validation of the results of this study and integration for multi-user interfaces [6, 7].

Figure 1. Limited area-based algorithm – depth-first search with disconnected procedures

References/Bibliography:

[1] ISO, “Systems and software engineering — systems and software quality requirements and evaluation (square) — measurement of quality in use,” International Organization for Standardization, Geneva, Switzerland, ISO ISO/IEC 25022:2016, 2016.

[2] K. H. Abbott, “Human factors engineering and flight desk design,” in Digital Avionics Handbook. CRC-Press, 1998, ch. 9.

[3] E. Suhir, “Human-in-the-loop (hitl): Probabilistic predictive modeling (ppm) of an aerospace mission/situation outcome,” Aerospace (Basel), vol. 1, 12 2014.

[4] M. Wang, C. Liu, and Y. Pan, “Computer-aided panel layout using a multicriteria heuristic algorithm,” The International Journal of Production Research, vol. 29, no. 6, pp. 1215–1233, 1991.

[5] G. Calvary and J. Coutaz, “Introduction to model-based user interfaces,” Group Note, vol. 7, p. W3C, 2014.

[6] Y. Zhang and Y. Sun, “Reuse of pilot motions for improving layout design of aircraft cockpit.” JCP, vol. 8, no. 9, pp. 2269–2276, 2013.

[7] G. Singer, “Methods for validating cockpit design the best tool for the task,” Ph.D. dissertation, Institutionen för flygteknik, 2002.

Achieving Consensus in Elixir using Raft

Matthew Alan Le Brun **Supervisor: Prof. Adrian Francalanza**

Course: B.Sc. (Hons.) Computing Science

Testing systems designed to run in a distributed-network setting may be described as difficult at best. Writing test cases for all possible scenarios for such an architecture is virtually impossible, due to the random nature of the computer network. One must consider message delays, drops and duplications, network partitions and worst of all, random combinations of these events. Run-time verification (RV) may instead be considered to determine the correctness of a software system within such an environment. RV refers to the tracing of events and process flow by means of a 'monitor' to determine whether the observed behaviour is complicit with defined invariants. This is how Graft – an Elixir implementation of the Raft consensus algorithm – was evaluated. Raft is an algorithm used to manage replicated logs, i.e. an event log that is duplicated across multiple servers within a computer network. It achieves this by orchestrating a cluster of machines to operate as a coherent group, capable of tolerating the failure of 'n' participants within a cluster of 2n + 1 machines. The orchestrator of the cluster is the leader, a server within the group elected within a time period to have full control over accepting client requests, replicating entries, applying entries when it is safe and replying back to the client when an answer has been computed. If a server is not a leader, then it is either a candidate (and waiting to receive majority votes to become leader) or a follower.

The main contribution of this project is Graft – an application that could be used to set up local or distributed

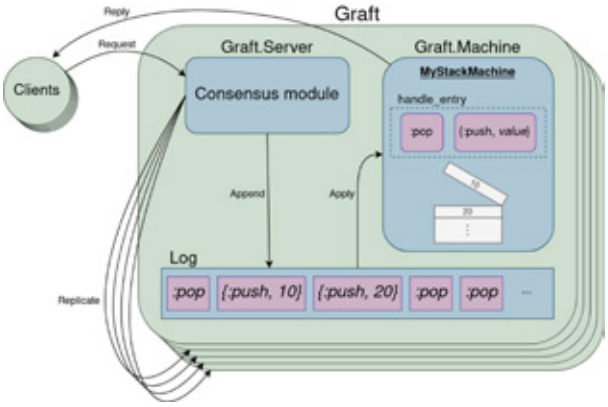

Figure 1. A visualised overview of the Graft application with a state machine designed to control a stack

Raft clusters with customisable state machines. Graft's correctness is supported through a variety of techniques, namely unit/integration testing and, previously mentioned, RV. The latter was achieved by using detectEr, which is an outline-monitoring tool for the Erlang ecosystem. Monitors were attached to servers within the Raft cluster to trace occurrences of events during the system's run-time. This ensured that the implementation adhered to the main invariants of the Raft algorithm. A total of four distinct monitor specifications were used to verify Graft, none of which found any violations for the invariants they represented.

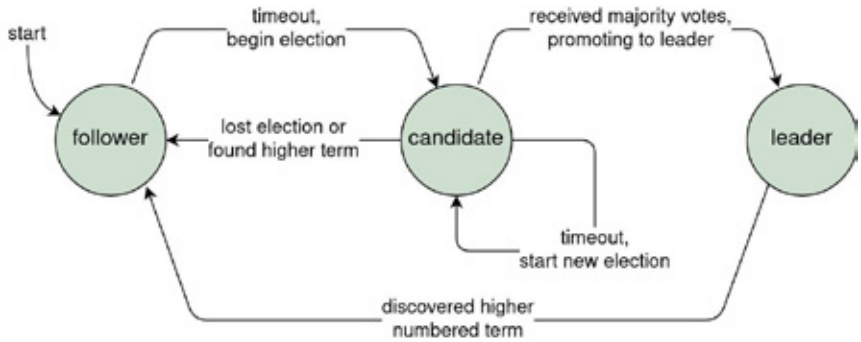

Figure 2. The conditions in which a server within the Raft cluster changes its state

Indexing of high-dimensional data in DBMS

Corinne Portelli **Supervisor: Dr Joseph G. Vella**

Course: B.Sc. IT (Hons.) Software Development

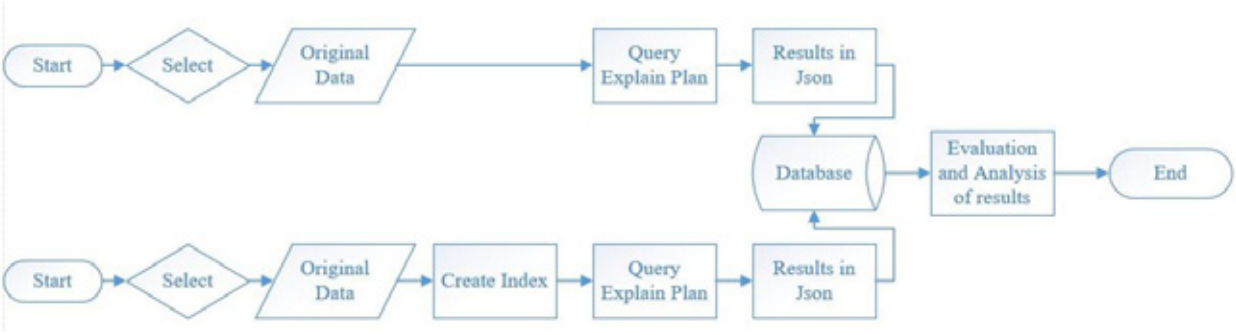

Figure 1. Logical-design diagram of the methodology used in the study

Novel applications rely on data with a high number of attributes and a much richer set of data types (e.g. spatial data). Furthermore, applications that gather data from sensors generate datasets that include numerous dimensions. It is, therefore, recommended that indexes to be utilised in database management systems (DBMS), in order to retrieve data more efficiently against the cost of maintaining them [1].

Many researchers have endeavoured to create indexing solutions for managing high-dimensional data. A staggering number of factors affect successful indexing on multi-dimensional data. Moreover, single-dimension indexing structure techniques do not generally apply in multi-dimensional databases.

When indexing spatial data, the R-tree indexing structure is normally used. However, when using indexes on PostGIS tables, PostGIS recommends the use of the GiST index framework. Furthermore, spatial indexing is provided with an R-tree over GiST scheme, as this framework allows data entries that are larger than 8Kb in size. This project plans to test whether a multi-dimensional index artefact would be worth introducing. Two types of datasets were used: a spatial dataset and a high-dimensional dataset. Different queries were run on the different datasets, and subsequently the query plan was recorded and analysed. The query planner was configured to run the queries under three different scenarios. The first scenario being with no indexes, a second scenario where the index was forcibly used, and finally giving a free rein to the enabled optimiser, i.e. where the query planner would be free to choose the execution plan.

This evaluation focuses on the differences of the recorded query plans, and calculations were made to determine when an index would be worth implementing, based on a breakeven analysis between query execution, index structure upkeep and use, and a simple amortisation of indexing-related costs (see Figure 1). Furthermore, this project aims to further reinforce the belief that indices should be used for high-dimensional data, and simultaneously providing a quantitative method to confirm whether to introduce these indices. One of the datasets used in the project was the publicly available spatial dataset of the public transport system in New York City. The other dataset that was utilised is a randomly generated table with a column containing a four-dimensional cube.

From the results obtained, one could conclude that the indexes were highly beneficial on large datasets. Also, when comparing the query plan chosen freely by the query planner with the manipulated query plans that were used, in most cases the query planner chose an execution plan similar to the one that was subjected to the involuntary index.

References/Bibliography:

[1] R. Bayer and E. McCreight, “Organization and maintenance of large ordered indexes,” Acta Informatica, vol. 1, p. 173–189, 1972.

Formal Model Extraction from Informal Descriptions

Jonathan Zammit **Supervisor: Prof. Ernest Cachia**

Course: B.Sc. IT (Hons.) Software Development

The translation of natural-language specifications to model-based specification could be defined as the translation from unstructured to structured-system specification. In this step, there is always the possibility of the presence of translation errors, such as human errors, which arise from unstructured text analysis [1,2]. This project aims to explore and study this translation process, with a view to propose an automated mapping paradigm to produce structured models from unstructured natural textual descriptions of system behaviour.

Figure 1. Example of natural language processing

After substantial research on both unstructured and structured languages, the project focused on the translation of natural language to Unified Modeling Language (UML), a graphical language that includes various non-mathematical software-modelling techniques, each aimed at visualising, specifying, constructing, and documenting software systems [3]. The transformation was carried out through two main steps: analysis and data extraction from natural language, and mapping the analysed text to a specific formalised model notation.

The first stage includes the synthetic reconstruction of the specification text and various natural-processing techniques. Sentence segmentation, tokenisation, lemmatisation, parts of speech tagging and named-entity recognition all played a part in the text-normalisation displayed in Figure 1. Through this process, the ambiguity, inconsistency and incompleteness of natural language are tackled, producing a standard format that was sufficiently adequate to act as an input for the mapping paradigm.

The mapping paradigm involves a series of rules and patterns directed towards identifying the components of the UML models. This project concentrates on the development of use-case diagrams, which are directed towards system-client interaction and highlighting system functionalities, and class diagrams, which provide visual schematics of the system that assist its design and implementation [3]. Elements and their relationships could be identified by using the first-stage output against the mapping paradigm. Through this step, a visual representation of the diagram is presented to the user.

The proposed fully automated system, as illustrated in Figure 2, provides the developer with a tool that could create diagrams with accuracy and ease, facilitated with optional manual refinement, thus being of help in the validation and verification of the system. This improves the system analysis by the client and the development team through the detection of errors prior to development. Thus, any risk could be better pre-empted and averted, making the development life cycle process more efficient, in terms of time and cost [4].

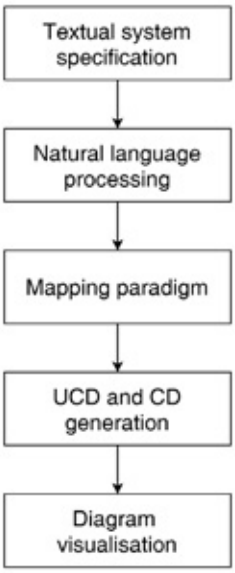

Figure 2. Flow diagram of the proposed solution

References/Bibliography:

[1] M. Leszek A and B. L. Liong, Practical Software Engineering. UK: Addison-Wesley, 2005.

[2] M. Muqeem and M. Rizwan, "Validation of Requirement Elicitation Framework using Machine," in IEEE International Conference on Control, Instrumentation, Communication and Computational Technologies (ICCICCT), 2014, pp. 1210–1216, doi: 10.1109/ICCICCT.2014.6993145.

[3] Unified Modeling Language, 2.5.1. Object Management Group, 2017.

[4] S. Kumar, R. Singh Suryavanshi, and G. Chandra, "Formal Methods:Techniques and Languages For Software Development," 2015.

Automated Face Reduction

Dejan Aquilina **Supervisor: Dr Ing. Reuben Farrugia**

Course: B.Sc. (Hons.) Computer Engineering

The General Data Protection Regulation (GDPR) [1] was implemented in May 2018, giving individual citizens the right to request the deletion of any personal data, including closed-circuit TV (CCTV) footage. Being generally installed in public spaces, such as shops, banks and in many streets, CCTV cameras are virtually ubiquitous. This makes the deletion of personal data from CCTV particularly problematic. Firstly, surveillance footage is required intact in the event that it might be called upon as forensic evidence in court. Secondly, obfuscating frames concerning a particular individual from a video would require substantial manual work, as the process would entail reviewing the entire footage, finding all the frames where the specific person or persons would be visible, and manually obfuscating them.

Semi-automated video redaction tools are currently available commercially. Among others, IKENA Forensic and Amped FIVE software packages [2, 3] allow the user to specify the area of interest to be obfuscated. With the use of semi-automated tracking techniques, the object of interest is followed throughout the footage. Although this tool facilitates the process, the identification of the person of interest would still necessitate the manual scanning of the video, which is a highly time-consuming task. Moreover, another major issue with the above-mentioned tools is that the licence to use them costs thousands of euros.

This final-year project involves the design and implementation of an open-source automated video redaction method that can automatically detect the face of the person or persons of interest in the video footage, and obfuscate them throughout the same footage. Deep learning was used to detect [4] and recognize [5] the object of interest, in preparation for face obfuscation. Dense optical flow [6] was subsequently used to track the face of the person of interest within the video, which was then obfuscated. This method achieved an accuracy higher than 90% in an indoor environment.

A video demo of the project can be found through the following link: https://youtu.be/1SeYgnt1Ny8

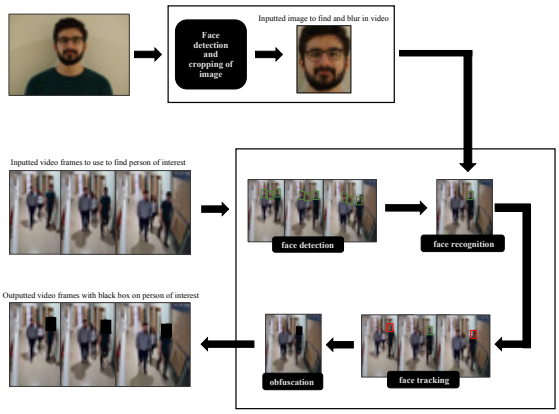

Figure 1. System diagram

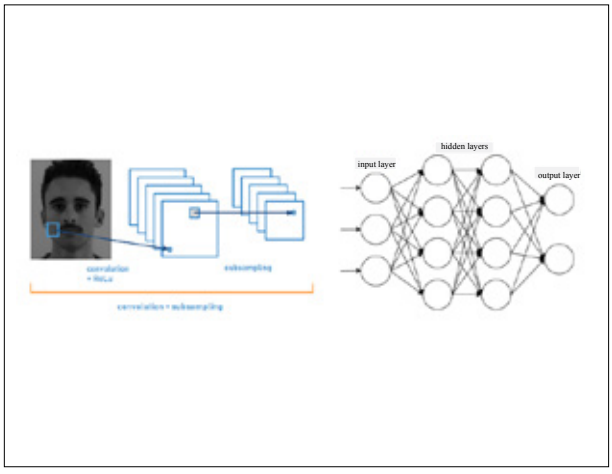

Figure 2. Convolutional neural network architecture

References/Bibliography:

[1] European Parliament, Council of the European Union. Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), 2016

[2] IKENA Software. “MotionDSP - Spotlight Software.” motiondsp.com. https://www.motiondsp.com/ (accessed May 6, 2020).

[3] Amped Software. “Amped FIVE | Forensic Video and Image Enhancement.” ampedsoftware.com. https://ampedsoftware.com/five (accessed May 6, 2020).

[4] J. Redmon and A. Farhadi. YOLO9000: Better, Faster, Stronger, 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), Honolulu, HI, 2017, pp. 6517-6525, doi: 10.1109/CVPR.2017.690.

[5] Omkar M. Parkhi, Andrea Vedaldi and Andrew Zisserman. Deep Face Recognition. In Xianghua Xie, Mark W. Jones, and Gary K. L. Tam, editors, Proceedings of the British Machine Vision Conference (BMVC), pages 41.1-41.12. BMVA Press, September 2015.

[6] Aslani, Sepehr & Homayoun Mahdavi-Nasab. Optical Flow Based Moving Object Detection and Tracking for Traffic Surveillance, 5th Iranian Conference on Electrical and Electronics Engineering, 2013.

Virtual Interfaces - Responsive and Adaptive (VI-RA)

Giulia Elena Caligari **Supervisor: Dr Vanessa Camilleri**

Co-Supervisor: Mr Foaad Haddod

Course: B.Sc. IT (Hons.) Artificial Intelligence

In Industry 4.0, the user is overwhelmed with all the information that is generated from the machines. This creates a problem, as informative irregular errors could get lost in the sea of data. Hence, virtual interfaces-responsive & adaptive (VI-RA) was designed employing anomaly detection to identify these abnormal entries within a multitude of data. VI-RA could then visualise data for each zone on site to the user through an augmented reality (AR) application, allowing for interaction with the display to reveal information about the errors.

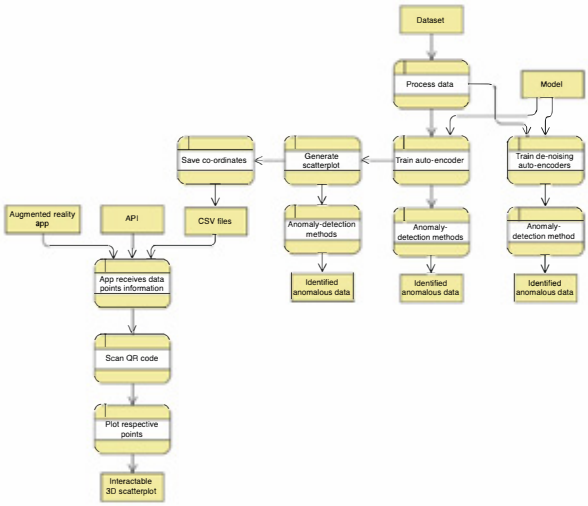

Figure 1. Data-flow diagram for VI-RA

VI-RA makes use of three main methods for the identification of anomalous data. The first method makes use of an auto-encoder that identifies an anomalous field by passing data through a trained model and indicating which entries return the largest error. The second method uses a de-noising auto-encoder to reconstruct missing fields within the dataset. This allows for a comparison between an expected and generated output towards singling out the anomalous field. The final function generates a scatter plot from the auto-encoder's thought vector so as to produce a visual representation of the dataset. The distance between the generated data points is then used to identify anomalous fields. Although this method managed to generate the desired results in this study, it proved to be the least consistent among the three. This was more evident as the complexity of the anomalies within the dataset increased. On the other hand, the auto-encoder and de-noising auto-encoder could return the desired anomalous entries with better frequency, thus proving to be the more promising methods for anomaly detection.

Figure 2. Diagram of the AR application, displaying the 3D scatterplot for a zone showing single-point data

An AR application was used to allow the user better visualisation and interaction with certain achieved results. The scatterplot co-ordinates for each zone within the dataset are sent over an API and received by the application, which displays a three-dimensional plot over a respective QR code for each zone. The points reveal information about their relative-error occurrence when interacted with, allowing for outliers to be easily identified, and the cause of their occurrence to be immediately evaluated.

Age Estimation using Deep Learning

Daniel Camilleri **Supervisor: Dr Ing. Reuben Farrugia**

Course: B.Sc. (Hons.) Computer Engineering

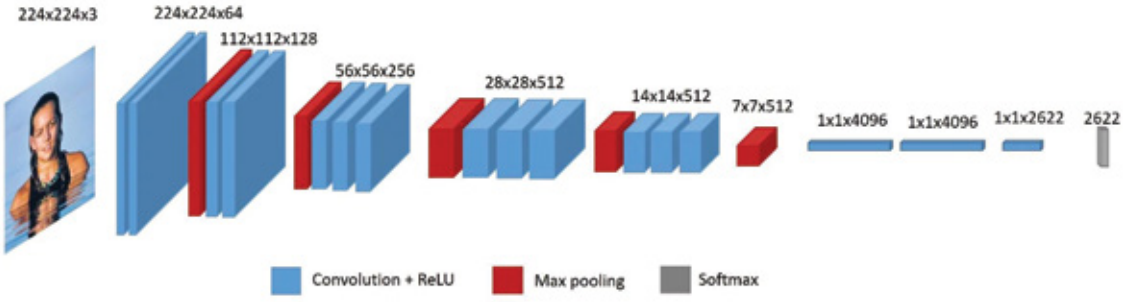

Figure 1. The age distribution of the IMDb-Wiki dataset

The progress of computer vision has made great strides in recent years, thanks to the significant research that was conducted exploring different applications of neural networks on images. One of these applications is age estimation.

In view of certain factors, such as face orientation, illumination quality, occluded faces, and black-and-white images – all of which could significantly affect the outcome – accurate age estimation could be difficult to achieve. However, recent advances in deep learning have achieved state-of-the-art performance on challenging datasets. The computer-vision software company Sighthound [1] claims to have achieved the lowest mean absolute error to date. In fact, the company also offers its neural network expertise as a service to retail businesses, enabling the latter to gather different statistical data about their customers, such as age and gender.

Age estimation is considered to be challenging in terms of regression analysis, due to the number of age-classification possibilities (ranging from 0-100+) and the extensive research conducted using different classification methodologies. Moreover it could be viewed as a texture pattern, in which the features could be used like local binary patterns, biologically inspired features and convolutional neural networks (CNNs) which, in recent years, experienced a surge of popularity due to the outstanding performance in face recognition.

Figure 2. Different examples of age estimation using the VGG-FACE transfer learning model

This project has investigated the use of deep learning for age estimation. Instead of creating a CNN and training it from scratch, this study draws from transfer learning, by using a pre-trained model called VGG-Face, which has been trained on thousands of images for face recognition. A support-vector regression (SVR) model was then trained on the features outputted by the CNN to reach the final age estimation, by fine-tuning the data obtained through VGG-face. The study made use of over 60,000 images to train and test the network, which achieved an accuracy of ±6 years. Although this number of images is not sufficient for training a neural network from scratch, it is adequate for the purpose of transfer learning.

References/Bibliography:

[1] https://www.sighthound.com/technology/face-detection/benchmarks/bao-and-afw
[2] IMDb-Wiki dataset https://data.vision.ee.ethz.ch/cvl/rrothe/imdb-wiki/

Workplace Assistant Augmented Reality

Gabriel Camilleri **Supervisor: Dr Vanessa Camilleri**

Course: B.Sc. IT (Hons.) Artificial Intelligence

Starting a job in an office tends to be very stressful for an intern or a new employee, and the first day at the office could be particularly challenging. It takes time for the new staff to adjust and to grow familiar with other employees and their duties, and how they could be of assistance to the existing staff. Furthermore, it could take some time for new staff members to learn the ropes and their purpose within the office building, while understanding and learning how to use certain equipment, such as an automatic key lock or even just a coffee machine. Therefore, the app being proposed – Workplace Assisted Augmented Reality seeks to identify the users' requirements for accomplishing a specific task within the workplace through user profiling and recommendation, whilst providing the relevant information for the users to learn and understand the environment around them via augmented reality (AR).

Figure 1. Model target-generation via Vuforia's deep learning techniques

The application is intended to guide new employees through an adapted process that would enable them to understand the environment around them, the equipment they might use daily, and directions for easily finding their way between offices. It is also intended for guiding the users, by providing them with relevant information in order to succeed in their assigned task. The system incorporates collaborative filtering and a similarity-based technique using the SVD++ model and item-to-item-based similarity respectively, to provide recommendations – along with deep learning and traditional computer-vision techniques using Vuforia – to provide augmented reality.

Using the above-mentioned tools, the system provides information about offices, directions towards specific offices, and information on how to utilise the coffee machine of the company, for which the application was tailor-made..

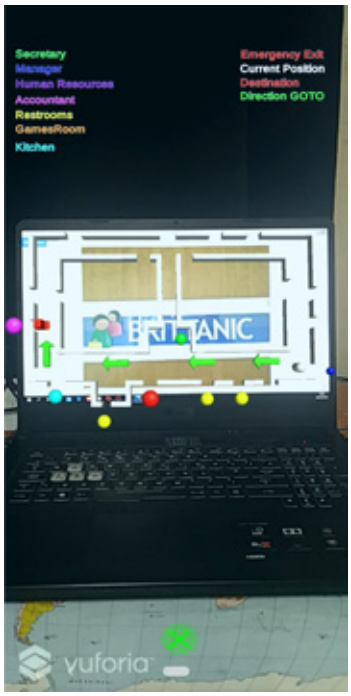

Figure 2. A 3D holographic map of the workplace, showing given directions towards the accountant's office, and recommended rooms and offices via collaborative and similarity-based techniques

From the tests performed, the indication is that the SVD++ model was the most efficient for the AR environment in comparison to other machine learning models. SVD++ achieved an average root-mean-square error and mean absolute error of 3.1226 and 2.6866, respectively. The AR component achieved promising distance, colour, rotation and occlusion variance values. User evaluation results have given positive indications that the AR application was usable, useful and effective. In addition, this study proposes a number of recommendations on how this application might be improved further.

A Study on the Effect of Target Object Size in Object Detector

Jacob Cassar Ellis | **Supervisor: Prof. Matthew Montebello**

Co-Supervisor: Mr Dylan Seychell

Course: B.Sc. IT (Hons.) Artificial Intelligence

Recent years have seen an impressive increase in the efficiency and accuracy of object-detection models through the use of region-based convolutional neural networks (RCNN). Many studies have been carried out that revolve around testing the limits of object detection. However, the detection of small objects still poses numerous issues. In fact, small-object detection is considered to be a significant issue for a number of reasons. One of these is the reduced accuracy of these systems resulting from relatively coarse feature maps, as they lack the information required for sufficient detection. Moreover, accuracy is compromised as these networks struggle to distinguish between foreground and complex backgrounds, such as rough terrain.

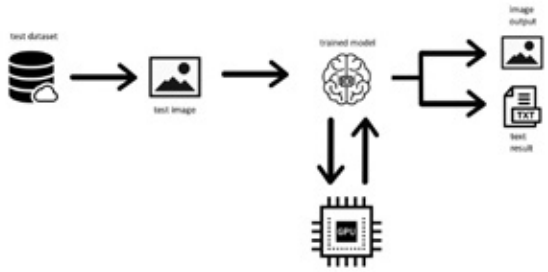

Figure 1. High-level architecture for testing an object-detection model

An experiment was undertaken to investigate the suitability of artificial intelligence (AI) techniques in analysing the effects on multiple-object detection models resulting from the distance between the capturing device and the object requiring detection. The set-up requires a custom dataset containing object instances that are identifiable by the implemented models, each with an object instance at a varying distance.

The starting point of the system is the sourcing of state-of-the-art object-detection models that are trained on the COCO dataset. To evaluate each model, the custom dataset was prepared in a private-environment set-up, with pre-established distance measurements and multiple small objects placed in plain backgrounds for detection. By changing the range of distances, each object instance underwent a change in spatial resolution in order to determine its effect on the accuracy of the model. Once the dataset was obtained, each of the models were implemented using a Python environment. Before applying the custom dataset, a series of baselines were initially tested to ensure that each model could correctly identify the selected object instance to be used as a ground truth. Once the baselines were established, each model was tested and evaluated using the custom dataset.

The designs and methods that were implemented were constructed in such a way as to provide a user-friendly system. Moreover, the aim of this project was to lay the foundations for future studies regarding small-object detection, by providing a different point of view to the issue of small-object detection.

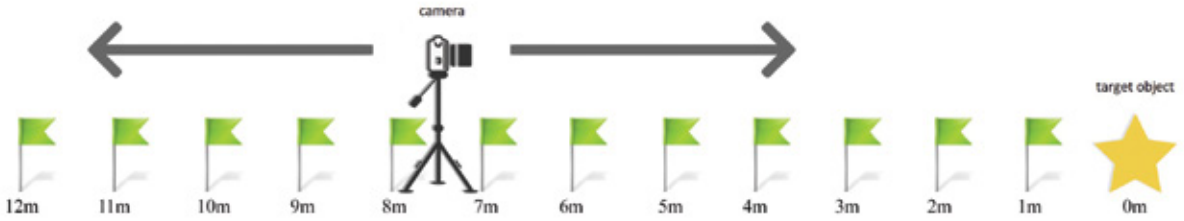

Figure 2. The architecture for a custom-dataset set-up

A data-analytic and machine learning approach to diabetes monitoring

Daniel Anthony Cilia | **Supervisor: Dr Michel Camilleri**

Co-Supervisor: Mr Joseph Bonello

Course: B.Sc. IT (Hons.) Software Development

Diabetes mellitus – commonly known as diabetes – is one of the most prevalent challenges in modern healthcare. Whilst presenting itself in a variety of types, its salient characteristic is the inability of the body to regulate blood sugar (glucose) levels, significantly hindering the individual's health and general well-being. Although a cure is yet to be found, diabetes could be kept in check through regular exercise and a well-balanced diet, coupled with keeping track of fluctuations of blood glucose through constant monitoring. The latter is considered by many as a necessary evil, as current monitoring solutions entail highly invasive procedures alongside other drawbacks that further tarnish the patient's quality of life [1]. Recent studies have presented interesting use-cases for machine learning (ML) algorithms to predict glucose levels within a certain time frame [2, 3].

The study was carried out within the framework of two main considerations. The first seeks to establish whether physiologic parameters, gathered from non-invasive sources, could be used to improve glucose predictive accuracy. The second consideration investigates whether the elimination of data gathered in an invasive manner would yield clinically acceptable results.

Multiple data analyses were conducted using time series data obtained from a sizeable ML dataset called OhioT1DM [4]. The first phase consisted of analysing and generating predictions with the aim of improving predictive performance. This was achieved by using feature-engineering techniques and splitting the dataset into different feature combinations. Data features were organised into the following groups: Glucose (G), Insulin pump (P), Fitness band (B) and Self-reported (S). Separate analytic steps (pipelines) were performed on each group, with the aim of producing a refined feature set, upon which ML models could be applied. Different combinations were tested on multiple linear regression (MLR) and XGBoost models. The results produced were evaluated to gauge the effectiveness of the input combination. In the second phase of experimentation, glucose-level values and their derived attributes were omitted completely from input

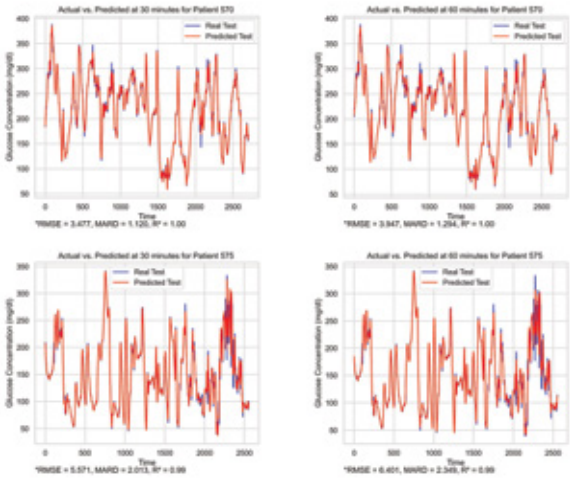

Figure 1. Prediction results from a multiple linear regression model trained on glucose values of up to two hours beforehand

features, and the corresponding predictive accuracy was evaluated. This was done to support, or otherwise, the ongoing research objective that non-invasive glucose monitoring could be achieved by means of multi-physiologic sensor monitoring. Root mean squared error (RMSE), mean absolute relative difference (MARD), R2 coefficient, and surveillance error grid (SEG) analysis were used as metrics to evaluate the produced results.

The findings suggest observable gains in predictive accuracy obtained using simple ML models, provided that the appropriate data-preparation mechanisms would be in place. This indicates that a simple and computationally lightweight model, such as MLR, could be used in, for instance, a mobile environment – with positive results. The importance of handling missing data is also highlighted, as features having significant gaps hindered predictive accuracy. In such cases, a totally non-invasive feature configuration would most likely yield poor results. For these reasons, further research in this field is warranted, particularly within the context of using more complex models to identify any potential hidden links among non-invasive features.

References/Bibliography:

[1] W. Villena Gonzales, A. Mobashsher, and A. Abbosh, "The Progress of Glucose Monitoring—A Review of Invasive to Minimally and Non-Invasive Techniques, Devices and Sensors," Sensors, vol. 19, no. 4, p. 800, Feb 2019." Sensors, vol. 19, no. 4, p. 800, Feb 2019. [Online]. Available: http://dx.doi.org/10.3390/s19040800" http://dx.doi.org/10.3390/s19040800

[2] F. L. Schwartz, C. R. Marling, and R. C. Bunescu, "The Promise and Perils of Wearable Physiological Sensors for Diabetes Management," Journal of Diabetes Science and Technology, vol. 12, no. 3, pp. 587–591, May 2018. [Online]. Available: https://doi.org/10.1177/1932296818763228" https://doi.org/10.1177/1932296818763228

[3] M. Gusev, L. Poposka, G. Spasevski, M. Kostoska, B. Koteska, M. Simjanoska, N. Ackovska, A. Stojmenski, J. Tasic, and J. Trontelj, "Noninvasive Glucose Measurement Using Machine Learning and Neural Network Methods and Correlation with Heart Rate Variability," Hindawi Journal of Sensors, vol. 2020, pp. 1–13, January 2020. [Online]. Available: https://doi.org/10.1155/2020/9628281

[4] C. Marling and R. Bunescu, "The OhioT1DM Dataset for Blood Glucose Level Prediction: Update 2020," Technical Report, 2020. [Online]. Available: http://smarthealth.cs.ohio.edu/bglp/OhioT1DM-dataset-paper.pdf

"To trust a LIAR": Does Machine Learning really classify fine-grained, fake news statements?

Mark Mifsud | **Supervisor: Dr. Colin Layfield**

Co-Supervisor: Dr Joel Azzopardi

Course: B.Sc. IT (Hons.) Software Development

Deep Learning

Fake News is a contemporary problem which causes serious social harm. People have been killed because of false assertions on online media [1] [2]. Early detection of fake news is therefore a critical problem in machine learning and NLP (Natural Language Processing); and a very challenging one.

This study attempted to automatically classify short claims, related to US politics, according to six levels of veracity, ranging from 'True' to "Pants on Fire" (ie. complete fabrications). These statements, taken from the LIAR dataset [3] were fact checked and pre-classified by experts for Politifact.com.

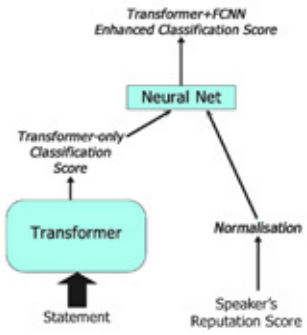

Figure 1. How the neural network used the speaker's reputation score and the classification output from the state- of-the-art accuracy score on fake news classification

In order to achieve a better accuracy score than previous studies, state-of-the-art, machine learning models known as transformers were used. Transformers are Deep Learning architectures meant to improve language understanding. Among these were Google's BERT [4]. BERT is the same language processing model that is used by the leading search engine to understand complex search queries. Other transformers like RoBERTa and ALBERT were also used. Such models have previously performed significantly well on several NLP tasks. To improve the classification of the claims, a simple neural network was also used to augment the transformer models so that they could utilise the claim maker's reputation score, thus enhancing the overall classification (Figure 1).

While the higher accuracy score was indeed achieved (Figure 2), the methods' ability to help mitigate the real-life problem of fake news' early detection, was still questionable. Flaws in the LIAR dataset were found, raising doubt about the validity of studies that used this dataset. For this reason, further experiments were done on the models built, to test their reliability. Flaws found in the models are discussed. Flaws include bias and the fact that they does not really model veracity which makes them prone to adversarial attacks.

Study	Model/Architecture	Data Used				Accuracy Score
		Statement	Metadata	Reputation	External	
Wang 2017	SVM	+				25.5%
	CNNs + LSTM	+				27.0%
Long et al 2017	LSTM + Attention	+				25.5%
This study	BERT_base	+				27.7%
	Roberta_Large	+				27.3%
	ALBERT_Large_V2	+				28.2%
Liu et al 2019	2 stage BERT_base+Attention	+				**34.5%**
Karimi et al 2018	CNN+LSTM	+				29.1%
Kirilin & Strube 2018	LSTM + Attention	+	+			41.5%
Pham 2018	Dual Attention	+	+			37.3%
	Memory Attention Network	+		+		44.2%
This study	BERT_base+NN	+		+		48.0%
	Roberta_Large+NN	+		+		47.9%
	ALBERT_Large_V2 +NN	+		+		**48.6%**
Wang 2017	CNNs + LSTM	+	+	+		27.4%
Long et al 2017	LSTM + Attention	+	+	+		41.5%
Kirilin & Strube 2018	LSTM + Attention	+	+	+		45.7%
Liu et al 2019	2 stage BERT_base+Attention	+	+	+		40.6%
Karimi et al 2018	CNN+LSTM	+	+	+	+	34.8%

Figure 2. A comparison of the accuracy score of previous studies against this study together with what deep learning model and data were used by each.

It was noticed that the statements in the dataset would require knowledge of the real world to accurately label as either true or false. The question inevitably arose about whether purely NLP-based, fake news classification, in general, can really be used to detect deception or whether it is an ill-posed problem.

A critique of this area of study was done. After scrutinising this study's own models, previous meta- studies and psychological studies about detecting deception, the author presents the argument that language-based, fake news classification should be treated as an ill-posed (unstable) problem. Some potential solutions to stabilise the problem in future studies are also suggested.

References/Bibliography:

[1] India's fake news problem is killing real people. https://asiatimes.com/2019/10/indias-fake-news- problem-is-killing-real-people/

[2] 'Hundreds dead' because of Covid-19 misinformation. https://www.bbc.com/news/world- 53755067

[3] William Yang Wang. 2017. "Liar, Liar Pants on Fire": A New Benchmark Dataset for Fake News Detection. ACL (2) 2017: 422-42.

[4] Devlin, J., Chang, M.W., Lee, K., Toutanova, K. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In: Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL-HLT).

[5] Ray Oshikawa, Jing Qian, William Yang Wang. 2018. A Survey on Natural Language Processing for Fake News Detection. ArXiv, abs/1811.00770.

[6] Chao Liu, Xinghua Wu, Min Yu, Gang Li, Jianguo Jiang, Weiqing Huang, Xiang Lu. 2019. A Two-Stage Model Based on BERT for Short Fake News Detection. In Douligeris C., Karagiannis D., Apostolou D. (eds) Knowledge Science, Engineering and Management. KSEM 2019. Lecture Notes in Computer Science, vol 11776. Springer, Cham.

Optimising station-stocking in bicycle-sharing systems

Julian Portelli **Supervisor: Dr Michel Camilleri**

Co-Supervisor: Dr Peter Albert Xuereb

Course: B.Sc. IT (Hons.) Software Development

This study aims to examine algorithmic approaches in increasing efficiency in bicycle-sharing systems (BSS), with the aim of increasing bicycle-sharing sales by ensuring availability when needed with the use of active restocking techniques. The study revolves around the dataset from a full year of BSS operation for the most popular bicycle-sharing service in New York City, Citi Bike.

Public bicycle schemes have received increasing attention in recent years. As of 2013, there were over 500 bicycle-sharing schemes around the world, and this figure is on a constant, albeit gradual, increase. More people are shifting towards renting bicycles for short trips due to the inexpensiveness of the scheme, as well as its convenience and minimal impact on the environment. Most of the current bicycle-sharing structures are automatic rental systems supported by information systems that can easily record data about a trip from the moment a user takes a bicycle from a pick-up point until the user deposits it at a docking station[1].

Despite being popular, the Citi Bike service is not always plain sailing. One of the drawbacks, for example, is that the high demand for this service by commuters heading for work results in a rapid exhaustion of bicycles at pick-up stations during peak morning hours in residential areas. Similarly, the docking stations in commercial areas fill up very quickly, rendering them unable to accommodate more bicycles [2].

Users of bicycle sharing systems unwittingly cause an imbalance in the system by creating demand in an asymmetric pattern, which may lead to unavailability of bicycles when the demand is high. This, in turn, brings about the potential loss of sales. In some situations, restocking a BSS (i.e. replenishing bicycle stations with full stock) after hours is not enough to warrant a well-balanced system.

This study focuses on the exploration of different means as to how to increase the efficiency of the scheme, primarily by rebalancing bicycle stations. At the most basic level, it would be necessary to determine the intensity of the demand experienced at the station(s) under observation at any given time. In this context, 'demand' refers to the number of bicycles taken out of a station at a given time. Thus, prediction of trip-demand at a station level is beneficial when seeking to predict the required stock at the respective stations. Knowing the possible number of future trips would be useful in the planning of when and how to restock during system operation. Hence, a scientific approach on this problem must be taken.

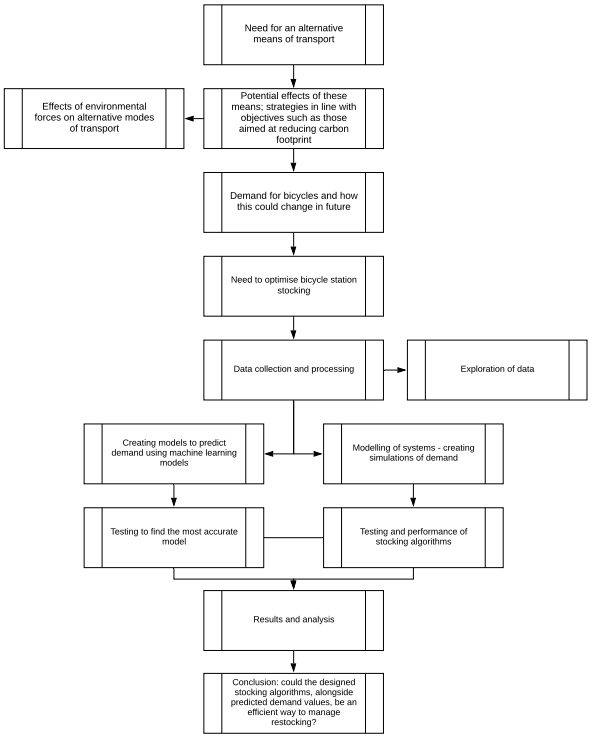

Figure 1. Relevance tree outlining the challenges of the bicycle-sharing process

Data exploration was used to identify patterns and relationships in trip usage. Furthermore, it was established that machine learning techniques would be necessary for prediction of hourly demand of trips per station at a particular time of day. In addition, this study presents a number of bicycle (re)stocking strategies in the form of algorithms, as well as an array of predictive models used to anticipate the aforementioned hourly demand.

The outcomes of the optimised best predictive model, combined with (re)stocking algorithms, would determine if system efficiency could be improved in terms of less opportunities lost in deposit and pickup This would optimistically result in better availability of bicycles in stations, thus possibly increase customer satisfaction, and in turn potentially translating into more sales.

References/Bibliography:

[1] W. Wang, "Forecasting Bike Rental Demand Using New York Citi Bike Data," Technological University Dublin, 2016.

[2] F. Chiariotti, C. Pielli, A. Zanella and M. Zorzi, "A Dynamic Approach to Rebalancing," Sensors, vol. 18, no. 2, p. 512, 8 2 2018.

A Citizen Approach for the Collection of Data to Train Deep Learning Models

Chantelle Saliba **Supervisor: Mr Dylan Seychell**

Co-Supervisor: Prof. Joseph Buhagiar

Course: B.Sc. IT (Hons.) Artificial Intelligence

Over recent years, technology in nature studies offered assistance in academic research and solutions for environmental crises. Machine learning is an approach that could continue to further the technological progress of this field. However, it is also widely acknowledged that machine learning techniques require a considerable amount of data, which could be particularly challenging to compile. Maltese botany is one area of interest that lacks the required quantity of data, resulting in limited technological advancements in the said field. Training a deep-learning network with insufficient data would cause overfitting. Therefore, auxiliary techniques would be required in order to overcome this issue.

The first part of this study investigates the training of a deep learning model that makes use of a limited training dataset utilising techniques such as data augmentation, data scraping and transfer learning. The deep learning model that was considered consisted of 50 categories, incorporating species endemic or indigenous to the Maltese Islands (such as the Maltese sow thistle). Data scraping was used for the collection of all available data. This technique did not generate sufficient training data, with sample sizes ranging between 10 and 217 for each category. Subsequently, data augmentation was used to enhance the dataset with augmented data. It was concluded that data augmentation performed well on both the training data and the testing data, generating a higher accuracy model than when data augmentation was performed solely on the training dataset. Different transfer learning models were also evaluated, and it was observed that the VGG-16 model outperformed the other models considered, generating an accuracy of 47.87%. This indicated that, although techniques such as transfer learning and data augmentation were implemented, the accuracy of the model was still relatively low.

While supporting the training of the deep learning model, these techniques were not a replacement to the data. This low accuracy of an improved off-the-shelf model showed the relevance of the initial hypothesis that citizen science (CS) would be required

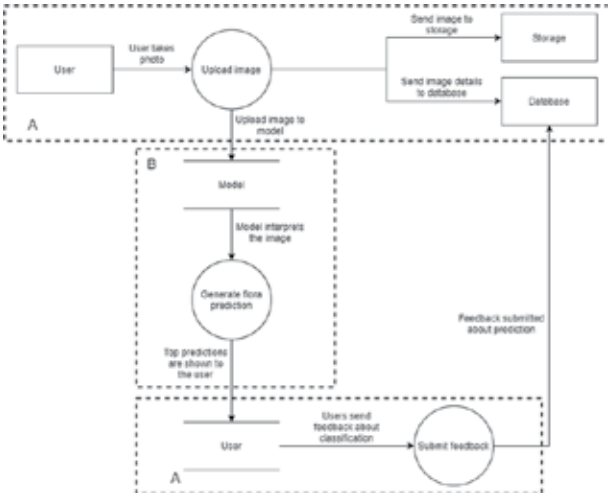

Figure 1. Data-flow diagram of the system

to enhance deep learning models. Consequently, this study has experimented with CS as a data-augmentation technique. CS and mobile-communication technology enable volunteering participants to contribute to the crowdsourcing of data. Since CS depends on the structure of society and culture, a questionnaire-based survey was conducted to determine the opinion of the general public regarding level of interest and willingness to participate. Out of the 243 respondents, it was only 13.2% who stated that they were not interested in such a mobile-communication system for crowdsourcing data. Due to the existing COVID-19 situation, the application could not be distributed and tested by the general public. However, the 230 images that had been gathered previously were used as an evaluation of the communication system. The collected data was also utilised to generate visualisations of the current distribution of Maltese flora.

Drawing from the research results, as outlined above, it was concluded that the use of CS would be essential for the improvement of deep learning models, when required to be employed in more widespread applications.

Automated Gait Analysis

Russell Sammut Bonnici | **Supervisor: Prof. Alexiei Dingli**

Co-Supervisor: Mr Dylan Seychell

Course: B.Sc. IT (Hons.) Artificial Intelligence

Gait analysis is the systematic study of walking patterns. Typically, expensive and intrusive marker- based methods, in conjunction with multiple infrared cameras, are used to produce kinematic data that concisely quantifies the walking behaviour of a person. Subsequently, specialists interpret conclusions from the kinematic data to make diagnoses. The aim of this research is to develop an automated alternative to marker-based methods for gait analysis.

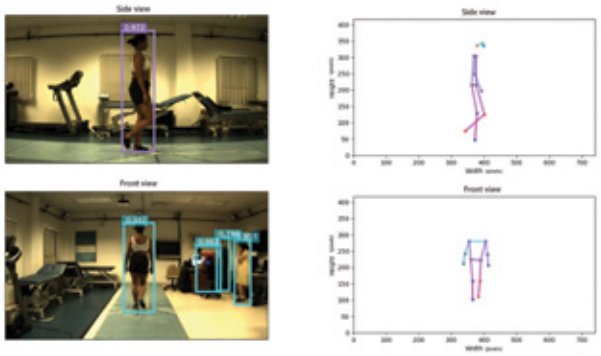

Figure 1. Pose estimation of a video frame with side and front views of the participant. Person detection is followed by keypoint detection resulting in pose-estimation graphs

The study proposes an automated method based on artificial intelligence that is able to achieve kinematic data consisting of varying left and right joint angles for hips and knees. All of this was obtained from videos capturing the side and front views of a walking subject. The method makes use of pose estimation, as well as a pipeline of techniques for calculating and processing kinematics. Included in the pipeline is a procedure based on ankle displacement for sampling walks into cyclic periods known as gait cycles.

Video data for the automated method, as well as marker data for a well-established marker-based method, was collected simultaneously in a biomechanics lab for comparison. Per participant, the automated method achieved a minimal error of about 3.14 degrees for left and right joint angles of hips and knees. This coincides with the measurement error of marker-based methods, which most studies have found to be less than 5 degrees (McGinley et al., 2009).

A binary classification experiment was performed with the kinematics of the automated method, in samples of gait cycles. Support vector machines achieved a 79% accuracy in detecting whether the walk of a person was normal or abnormal. The automated method has proven to be more suitable for data collection and processing than marker-based methods, and it may be utilised for further gait-classification procedures. It reduces the effort and financial investment needed for gait analysis, leading to a wider diffusion in the healthcare and research communities. The proposed automated method could also serve as a basis to train a classifier with sufficient kinematic data to accurately deduce gait-related conditions, and minimise the involvement of specialists for diagnoses.

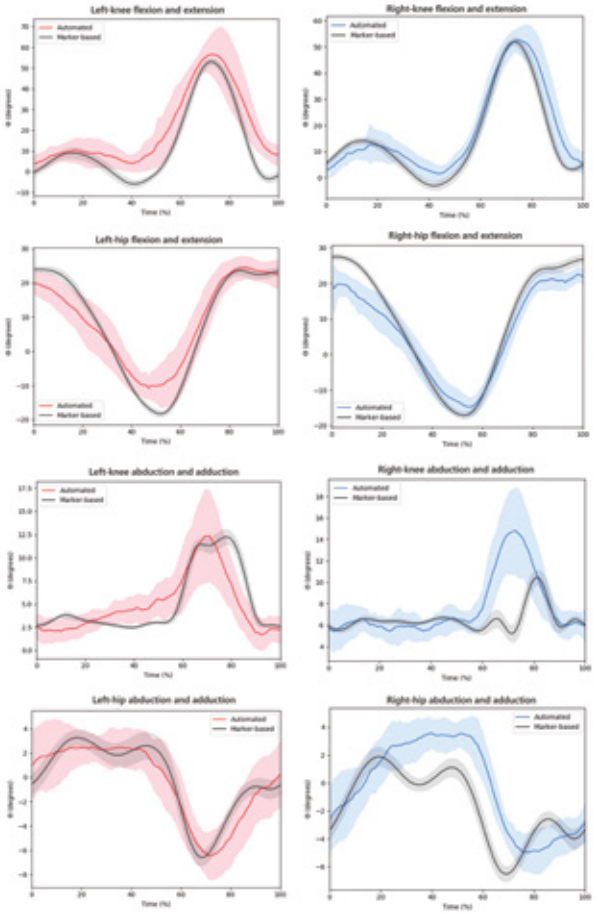

Figure 2. The average and standard deviation of kinematic signals of the proposed automated method (in colour), superimposed on the lab's marker-based method (in grey). This shows a sufficient similarity for effective and automated gait analysis

References/Bibliography:

[1] Jennifer L. McGinley, Richard Baker, Rory Wolfe, Meg E. Morris, The reliability of three-dimensional kinematic gait measurements: A systematic review, Gait & Posture, Volume 29, Issue 3, 2009, Pages 360-369

Teaching AI through Augmented Reality

Francesca Silvio **Supervisor: Prof. Matthew Montebello**

Co-Supervisor: Dr Vanessa Camilleri

Course: B.Sc. IT (Hons.) Artificial Intelligence

Artificial intelligence (AI) has pervasively penetrated every aspect of our life, impacting the way society operates in general [1]. Such an emphasis on this technology highlights the importance that the general public could clearly comprehend and appreciate the nature and scope of AI. This awareness must be nurtured from a young age, as children today constitute tomorrow's society, and are therefore required to be fully conversant with AI and how it could render their lives easier [2], [3].

This final-year project seeks to address the challenge of teaching AI concepts to children, by offering a digital tool to help teachers/educators achieve this. This digital tool is an Android application that was developed using Unity which in turn uses the Vuforia Augmented Reality software development kit. The application also makes use of the Google Cloud Platform (GCP) and Google App Engine in order to expose a Python-based web Application Programming Interface (API) containing convolutional neural networks (CNN) developed using Keras. The application is in the form of a game in which children need to teach an augmented reality (AR) robot how to classify different colours by providing images to train the CNN. The application provides step-by-step instructions on what needs to be done by the user as well as simple explanations on the machine learning (ML) concepts being used. These concepts include: neural networks, training a neural network, and using neural networks for predictions and classifications. The gamified system is intended to encourage children to take interest in what they are being taught [4].

Whereas the interface of the application was designed using Unity, the CNN was designed using Python. The study discusses how the application transfers data between the Unity interface which uses the C# Programming language and the CNN, written in Python, through the means of a Google Web API. The expected outcome of this final-year project is to provide an entertaining teaching method to motivate children to learn important concepts of ML, this being a very important area of AI. The proposed method has been deemed to be more engaging, thus rendering the topic easier to understand than if it were taught through traditional teaching methods [5].

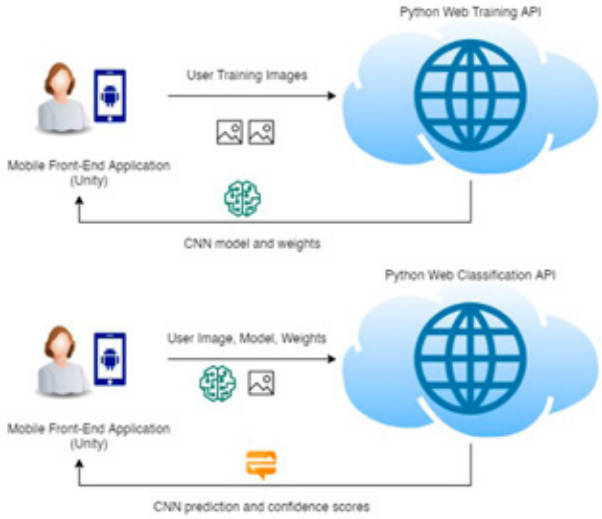

Figure 1. Flow of data between the application interface and the cloud-based API.

The effectiveness of the application was evaluated by comparing the level of knowledge gained after watching a video and using the application. The video was created specifically for the purpose of the study, and designed to convey the same information as the application. Target users (ie. children) were asked to answer the same questions after watching the video and using the application. The quality of these answers has been reviewed in order to identify whether the participants gave better answers after having used the application.

References/Bibliography:

[1] C. Baciu, D. Opre, and S. Riley, "A new way of thinking in the era of virtual reality and artificial intelligence," 06 2016.

[2] A. Sangra and M. Gonzalez-Sanmamed, "The role of information and communication technologies in improving teaching and learning processes in primary and secondary schools," ALT-J, vol. 18, no. 3, pp. 207–220, 2010.

[3] D. Kawade and S. Kulkarni, "Use of ICT in primary school," Pioneer Journal, 2012.

[4] J. Lee and J. Hammer, "Gamification in education: What, How, Why bother?," Academic Exchange Quarterly, vol. 15, pp. 1–5, 01 2011.

[5] A. Stott and C. Neustaedter, "Analysis of gamification in education," Surrey, BC, Canada, vol. 8, p. 36, 2013.

GAImE – Investigating Game AI for Enhanced User Experience

David Vella **Supervisor: Dr Vanessa Camilleri**

Co-Supervisor: Prof. Vince Briffa

Course: B.Sc. IT (Hons.) Artificial Intelligence

Over the last few decades, the business of digital and video games has risen exponentially. Game artificial intelligence (AI)in particular is a niche that has been gathering considerable momentum within the game-development industry [or sector]. This is mainly due to its higher-fidelity graphics and more engaging game plots and narratives than those present in 'standard electronic games. Research and development in the area is registering increased progress in a number of AI technologies.

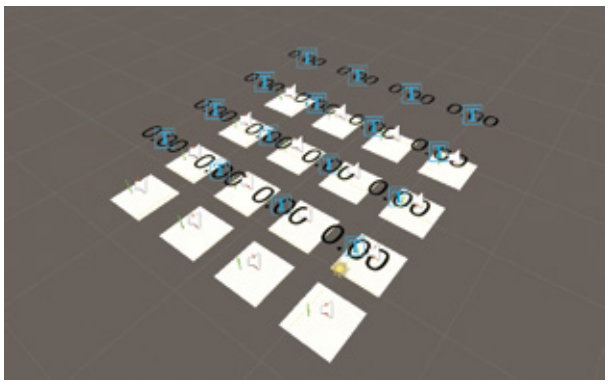

Figure 1. Unity training area for reinforcement-learning agents

Game AI research, which has been ongoing over the past 20 years, has also seen an increase in the investigation of how AI could be used to enhance the user gameplay, experience, rather than merely playing to win against human players. The concept in applying AI techniques is to equip non-player characters (NPC) with intelligence through the use of different techniques. NPCs are AI-controlled characters within the video game, which characters would have different typologies, such as companions and enemies.

This study investigates different AI techniques for the generation of NPCs, and how these might in turn affect the user's experience, by evaluating user engagement and satisfaction. Game studies indicate that the use of attractive graphics, good sound-quality and music, are essential game elements in improving user satisfaction. However, this study proposes an additional element to further ensure user satisfaction in gameplay, namely intelligence in NPCs. As a proof of concept, a survival game was developed using the cross-platform game engine, Unity 3D. Different AI techniques, such as finite-state machine (FSM), behaviour-based NPCs and reinforcement-learning (RL) agents were applied to the game's NPCs, and these were in turn tested and evaluated from the user-experience perspective. Whereas, in the case of behaviour-based NPCs the actions were predefined, reinforcement-learning agents learned the actions depending on the in-game rewards obtained.

Figure 2. A scene from Game World, which was developed specifically for this study

The satisfaction of the user was calculated by using the Game User Experience Satisfaction Scale (GUESS). This tool offers a good indication of the aspects of the game about which the user would feel more satisfied. The results obtained revealed a total average satisfaction of 5.8/7 on the GUESS scale. The survey results also indicated that users preferred the behaviour-based approach in favour of reinforcement learning. One possible explanation is that this particular group of users preferred to predict and anticipate the actions and reactions of the behaviour-based companion NPC, since all of its actions were predefined. It is to be noted that these preferences are subjective, depending on the type of audience playing the game.

Named-entity recognition for Maltese: A scenario for a low-resource language

Deborah Vella **Supervisor: Dr Claudia Borg**

Course: B.Sc. IT (Hons.) Artificial Intelligence

Named-entity recognition (NER) is a subtask in the field of natural language processing (NLP), whereby named entities such as 'person', 'organisation' and 'location' are identified and labelled in text. NER is a substantial contribution to information extraction, as it identifies the named entities from which the required information could be extracted, such as how they relate to each other. This is also known as 'entity linking'.

In general, state-of-the art NERs are trained on large corpora, with named entities already tagged through human-annotation initiative. However, not all languages have the benefit of a sufficiently large corpus. In fact, Maltese lacks both available NER-annotated datasets and previously created Maltese NER models. Hence, the aim of this study was to conduct research about previous low-resource NERs to obtain enough knowledge to contribute towards resolving the challenging task of creating and evaluating the first NER system for Maltese.

In this project, a small dataset of 500 sentences extracted from the MLRS corpus was created. Sentences are manually annotated at word level using the following categories: 'Person', 'Location', 'Organisation' and 'Miscellaneous'. These are further annotated using the BIO tagging system which indicates whether the words are found in the (B)eginning, (I)nside or (O)utside of a named entity. In order to broaden the dataset, transfer learning was tentatively applied by including datasets from other languages, namely English, Italian, Spanish and Dutch.

The experiments sought to evaluate the use of two techniques – Conditional Random Fields (CRF) and bidirectional long short-term memory conditional random fields (BiLSTM-CRF) as a deep learning approach. These experiments also required considering a number of scenarios, since there were no specific annotation guidelines for Maltese. Initially, tags were limited to the 'Person', 'Organisation' and 'Location' labels, with a later introduction of the 'Miscellaneous' tag for further experimentation. It was also necessary that analysis of the tags would match what was available in the selected multilingual NER datasets to streamline the transfer learning for the Maltese annotations. We also experimented with the size of the multilingual corpora to analyse the impact that other languages could have

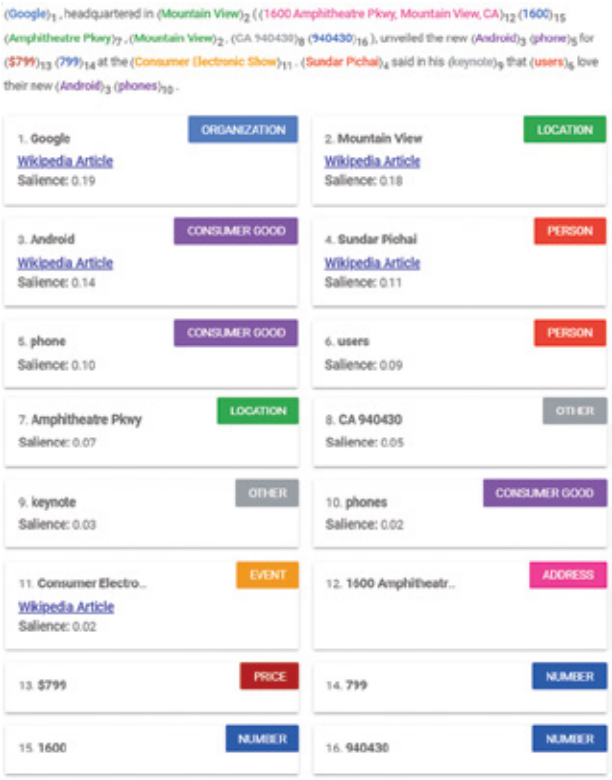

Figure 1. A visualisation of the typical output of an NER system, which has been extracted from the Google Cloud NER API demo output

on Maltese NER system. This was done incrementally with the first corpus containing Maltese only, and then the others containing one of the following amounts of sentences from each language: 200, 300, 400 and 500.

This project has sought to demonstrate the feasibility of transfer learning in training an NER system for Maltese in a low-resource setting. The experiments resulted in a large number of setups, totalling 40 distinct experiments. The best results were obtained by three equally successful systems achieved by the BiLSTM-CRF experiments. One of these systems was trained on Maltese and 300 extra sentences from the other languages without making use of the 'Miscellaneous' tag. The other two systems were trained on Maltese, together with 400 and 500 extra sentences from all the other languages, excluding Dutch.

Identifying bottlenecks in operational workflows at the Accident & Emergency Department at Mater Dei Hospital

Marco Attard **Supervisor: Mr Joseph Bonello**

Co-Supervisor: Dr Michel Camilleri

Course: B.Sc. IT (Hons.) Computing and Business

Operational research is an important subfield of applied mathematics that provides decision-makers with quantitative information derived from mathematical modelling and statistical analysis. The earliest applications of modern-operations research were to aid decision-making in World War II manoeuvres. Today, it is frequently used over a wide array of industries to analyse and optimise complex systems.

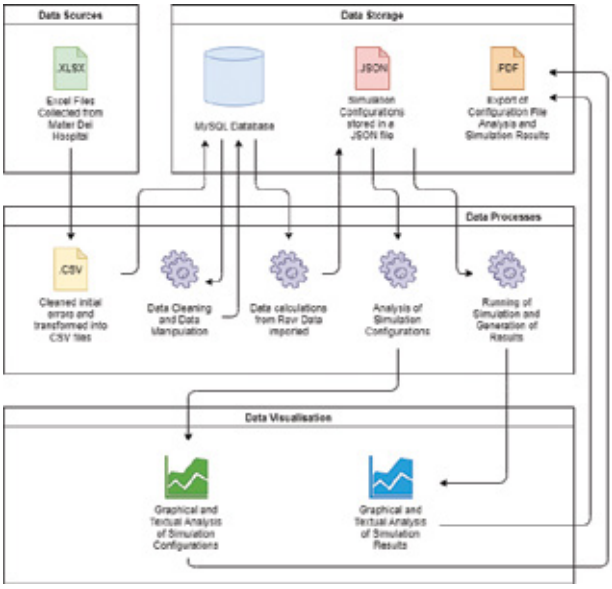

Figure 1. The data flow and processes

The two major subdisciplines of operational research followed throughout this study were: simulations and stochastic modelling. Computer simulations are models that represent real-life processes and systems, developed in order to analyse how they work. Additionally, these representative models are used to run calculations and make predictions. A stochastic simulation is a special type of simulation, which has variables that change randomly according to the stated probabilities.

In this study, a stochastic model of the operating procedures of the Accident & Emergency (A&E) Department at Mater Dei Hospital (MDH) was developed to represent the system under study. This representative model was developed through a dataset acquired from the said A&E Department. Extensive data-cleaning was performed on the dataset prior to developing the simulation tool. Once this data-manipulation process was completed, the current configuration of the operating procedures was derived and stored in a JavaScript Object Notation (JSON) file. This JSON file was then used in the main process of the simulation.

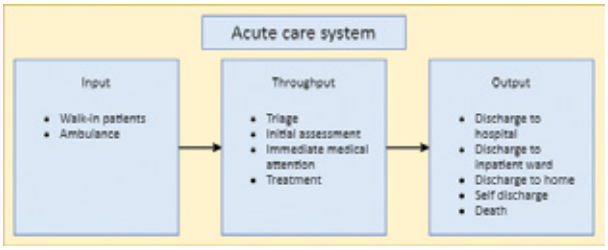

Figure 2. The input/throughput/output model

The developed tool simulates the input/ throughput/ output model applied to individual patients upon admission to the A&E Department at MDH. Through the use of the pre-calculated probability distributions, the tool determines how many patients are admitted each hour, the duration of each step of their stay and the sequence of steps they follow from admission to discharge.

The purpose of this research is not just to detect flows that could exceed the A&E key performance indicator (KPI) for the discharging of patients, but also to provide medical professionals with a configurable simulation tool, within which proposed changes to the operating procedures of an A&E department could be tested and evaluated. This tool seeks to enable stakeholders to successfully implement the proposed changes to the operating procedures after analysing and verifying their effects, thereby safeguarding the patients' wellbeing and improving the quality of the healthcare provided.

Detecting comparatives in images

Lara Brockdorff **Supervisor: Prof. Ing. Adrian Muscat**

Course: B.Sc. (Hons.) Computing Science

Image understanding (IU) is a subfield of computer vision (CV), which seeks to detect the semantics of an image. Object detection in images is the first step in IU. However, a list of objects alone is not enough for most applications in a number of areas, such as robotics, image-description generation and visual question-answering. Previous work dedicates significant effort in the detection of attributes (e.g. green box) and relations (e.g. on top of). On the other hand, comparatives (e.g. larger or taller than) is a relatively unexplored area. In this project In the absence of an existing dataset, this project required the building of data-driven pattern recognition (machine learning) models are built to detect comparatives in images. Moreover, a literature survey on gradables and comparatives was carried out before collating a suitable dataset.

Figure 1. Example of an image from the collated dataset

In the absence of an existing dedicated dataset, one was collated from readily available, human-annotated, easily obtained datasets, more specifically and primarily from: single-label annotation datasets; image-description datasets and visual question-answering datasets. The

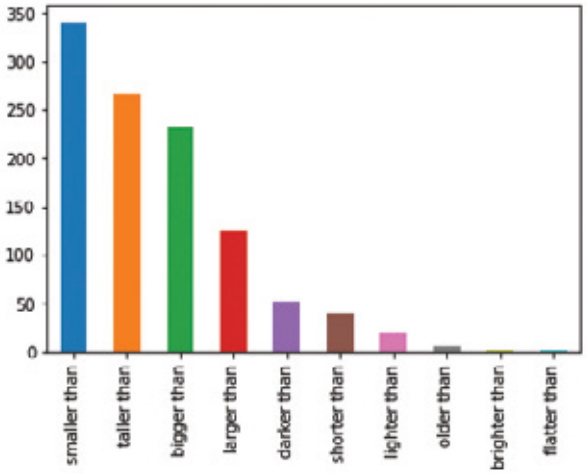

Figure 2. Distribution of classes in the collated dataset

dataset consists of images and identified object pairs with respective bounding boxes, and the corresponding relevant comparative, used as a gold label for the models to be built. This could be seen in Figure 1, where the comparative 'bigger than' was assigned to the object pair ('teddy bear', 'person'), and the objects identified with bounding boxes (blue and red respectively). The models make use of language features – as well as geometric features computed from the selected bounding boxes – in order to predict the comparatives relating to the identified pair of objects. The models developed were analysed and compared.

This study also discusses the challenges that arise in the tasks that attempt to extract semantic value from images using a set of features, as well as the challenges due to the ambiguity of the comparatives from a language perspective. The list and distribution of comparatives from the collated dataset can be seen in Figure 2.

The visual comparatives dataset and the results and critique of the models developed for the purpose, set the foundations for future work in this area.

Reducing risk in road network traversal using past road accident data

Miriana Buhagiar **Supervisor: Dr Peter Albert Xuereb**

Co-Supervisor: Mr Joseph Bonello

Course: B.Sc. IT (Hons.) Computing and Business

The World Health Organisation estimates that as a result of road traffic accidents, the lives of about 1.35 million people are ended each year [1]. Traffic accidents lead to loss of life and property; therefore, it is imperative to understand road data to provide safe traffic movements to the road users and ensure their safety [2].

The main objective of this study is to contribute towards creating safer roads by generating awareness amongst drivers of the risk of using particular routes. It is difficult to stop road accidents from occurring especially with the increase in vehicles on the road that we have seen in recent years. However, it is not impossible to reduce the impact of road accidents on drivers by extracting knowledge and patterns from reported data and informing drivers of potential risks.

This research project presents a probabilistic approach on how to measure and determine the risk of dangerous roadways. Data on road accidents from 2015 to 2019 from two separate data sources, the police and local wardens, was obtained and cleaned, processed and investigated to create a clear picture of the unsafe and hazardous roads in Malta. This work presents a methodology for identifying crash hotspots in terms of injury cost, and identifying high-risk locations by using the frequency of reported accidents. Every road that had any collision from was assigned a 'risk factor' from one to five: the higher the factor, the greater the amount of reported accidents in that particular road. On the other hand, every road where an accident that involved one or more injuries was reported was given a colour; either amber, red or black depending on the number of people injured when compared to those not injured in accidents in that road. Those roads labelled as an amber have a low probability of injury cost and loss when compared to the red and black roads, that are progressively worse.

An interactive map was built so that drivers can identify risky and hazardous roads when planning their driving trip or during their travel by car. The driver chooses the time of day (Morning, Afternoon or Evening), during which s/he intends to perform the journey and inputs the start and destination points of the trip. The route between the two points will be colour coded, depending on the risk factor of that road at that chosen time of the journey. If the route passes through injury-prone roads, these are marked on the map.

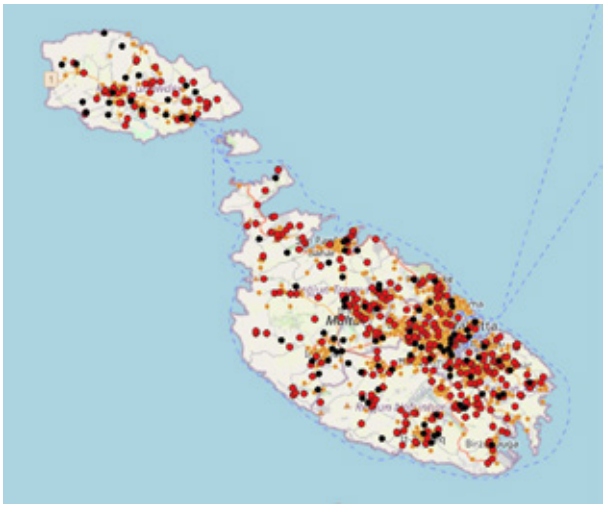

Figure 1. Accident hotspots in Malta and Gozo, colour-coded according to level of risk

The research findings are visualised on an interactive map, and the resultant probabilities show that Malta has on average more frequent collisions than Gozo. However, when taking injuries into consideration, Gozo has a higher number of percentages of red and black injury roads than Malta. Additionally, the roads that have the highest number of reported accidents, are not the most dangerous in terms of the injuries sustained during impact. Furthermore, it could be concluded that the roads with the highest number of accidents experience heavy traffic.

Therefore, traffic congestion leads to more accidents although, drivers are not usually injured.

References/Bibliography:

[1] World Health Organisation, Global status report on road safety: time for action, Geneva: www.who.int/violence_injury_prevention/road_safety_status/2009, 2009.

[2] F.Malik, S. Jabber and I.Rashid, "Road Accidents and Prevention," International journal of egineering development and research, no. 5, pp. 40-46, 2017.

Voynich - Hoax?

Adriana Camilleri **Supervisor: Dr Colin Layfield**

Co-Supervisor: Dr Lonneke van der Plas

Course: B.Sc. IT (Hons.) Software Development

The Voynich Manuscript – considered to be "the most mysterious manuscript in the world" [1] [2] – has been subject to intense efforts to decode its mysterious text. To date, all such attempts have been unsuccessful.

The manuscript is believed to be divided into six major categories, and this supposition is based on the illustrations on the pages that act as indicators of the presumed topics. These have been organised into the following categories: herbal, astronomical, biological, cosmological, pharmaceutical, and recipes [1].

Figure 1. A sample page from the original manuscript

The goal of this research is to attempt to examine whether the Voynich Manuscript is an elaborate hoax or indeed an unknown language. If the manuscript is the 'real deal', – as suggested by Reddy and Knight [3] that it is written in a real, but unknown language – then, it would be legitimate to surmise that, as a rule, the text within the respective categories is more semantically related within a given category, than is the text across different categories.

To test the above hypothesis, this study adopted a statistical approach in analysing the manuscript. The research has been performed using the Extensible Voynich Alphabet (EVA) – a transcription system that is digitally available to the public. Figure 2 is a snippet of the dataset using EVA and is a transliteration of the original text, a sample of which is shown in Figure 1. Since this dataset is an ongoing work and contains terms that have more than one interpretative reading, pre-processing techniques have been adopted to minimise any ambiguities. Various machine learning algorithms, including classification and clustering tools such as: support-vector machine (SVM), k-nearest neighbours (kNN), multinomial naïve Bayes, k-means and No- k-means have been utilised. The most reliable results were achieved when text from each category was extracted by page rather than on a line-by-line basis. This is in line with expectations, because a page of text offers a wider context than a line, and when seeking to assign a page or a line to a specific topic, the former would be easier to assign, since more words would be present and can be related to a specific topic [4]. Through different metrics and diagrams, these algorithms show clear and expressive results. Similar distributions over lexical items were found among the categories, and less so between the different categories, which provides evidence of an underlying lexico-semantic categorisation that is in line with the categorisation indicated by the illustrations. This contributes towards supporting the hypothesis of the existence of a genuine message within the manuscript.

Figure 2. Dataset snippet based on an EVA transcription of the sample page in Figure 1

References/Bibliography:

[1] R. Clemens and D. Harkness, The Voynich Manuscript, New Haven, Conn./London: Yale University Press, 2016.

[2] S. Skinner, R. T. Prinke and R. Zandbergen, The Voynich Manuscript: The World's Most Mysterious and Esoteric Codex, Watkins Publishing, 2017.

[3] S. Reddy and K. Knight, "What we know about the Voynich manuscript," in In Proceedings of the 5th ACL-HLT Workshop on Language Technology for Cultural Heritage, Social Sciences, and Humanities, Chicago, 2011.

[4] C. C. Aggarwal

3D Facial Reconstruction from 2D Portrait Imagery

Matthew Caruana **Supervisor: Dr Joseph Vella**

Course: B.Sc. IT (Hons.) Software Development

Facial reconstruction is a highly active field of research, and is applied frequently in the areas of facial rehabilitation [1] and forensic facial reconstruction [2]. Academic efforts in the field have been geared towards improving the accuracy and efficiency of the 3D models, which are generated from the input 2D images containing distinguishable human facial features with enough intra-variety. Such images allow researchers to generalise an average layout of features that is attributed to a generic human face. However, research in this field is invariably having to contend with the limitations inherent to 2D images, which slow down research progress. These limitations are largely due to: occlusions, the restricted variety in the features of a human face, different facial poses, and even partial views.

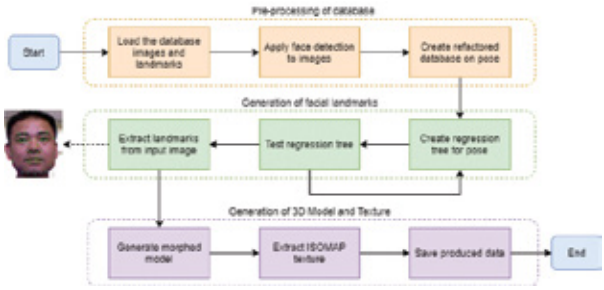

Figure 1. Proposed system pipeline

This project offers a system, outlined in Figure 1, consisting of the creation of specific regression trees for homogenizing human faces, and then applying the regressed face to an input 2D image to be reconstructed in 3D. These trees contain placements of a pre-defined set of landmarks that are found within the input 2D face, where applicable. A 3D morphable model is a collection of 3D surfaces, in which all vertices in every surface map correspond to the vertices of the different surfaces found within that same morphable model [3]. These points are then applied to the 3D morphable model which, after a series of transformations, outputs a reconstructed 3D model that would be adjusted to the specific features found in the input image.

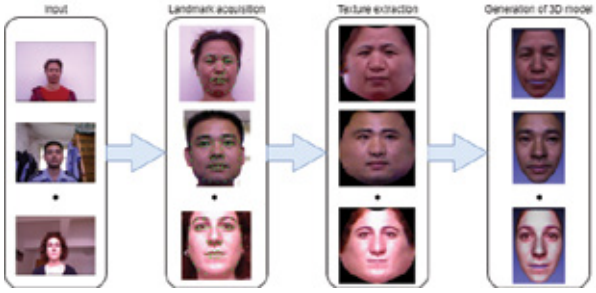

Figure 2. System flow results

The evaluation process used here involves determining mathematically how one could compare different models from the same person. This was achieved through the use of mesh-analysis metrics such as the 75th percentile, root mean square and geometric mean [5]. The result is a comprehensive 3D reconstruction system as shown in Figure 2, which generates a 3D model with its corresponding texture through landmark extraction from a single 2D portrait image.

References/Bibliography:

[1] G. Storey, R. Jiang and A. Bouridane, "Role for 2D image generated 3D face models in the rehabilitation of facial palsy," Healthcare Technology Letters, vol. 4, (4), pp. 145-148, 2017.

[2] Y. Hu et al, "Automatic 3D reconstruction for face recognition," in Sixth IEEE International Conference on Automatic Face and Gesture Recognition, 2004. Proceedings. 2004, pp. 843-848.

[3] V. Blanz and T. Vetter, "A morphable model for the synthesis of 3D faces," in Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques, 1999.

[4] C. Cao et al, "FaceWarehouse: A 3D Facial Expression Database for Visual Computing," IEEE Transactions on Visualization and Computer Graphics, vol. 20, (3),2014, pp. 413-425.

[5] F. Zuzana, "Comparison and Analysis of Multiple 3D Shapes", MSc, Masarykova Univerzita, Brno, Czech Republic. 2015.

Analysing News Portal Comments

Mikhael Cutajar **Supervisor: Dr Charlie Abela**

Course: B.Sc. IT (Hons.) Artificial Intelligence

The popularity of online news portals over recent years has resulted in these repositories of articles being augmented with readers' comments in response to the respective articles [6]. The main objective behind this project is to investigate the relation between articles and commentators, and to analyse and visually represent how different users comment across articles. This information includes data, such as the manner in which each user reacted to different articles, and which topics they preferred commenting about. An overview of the main components, which will be discussed in further detail, has been provided in the design flowchart accompanying this abstract.

The first task was to obtain access to news articles of a popular local news portal and store these in a MongoDB database. The extracted data covered a period of 15 months, ranging from January 2019 to March 2020, and contained 14,188 articles and 228,249 comments from 5385 unique users.

One of the main challenges encountered was to detect the topic/s associated with the articles. The initial attempts revolved around supervised learning techniques that required using a labelled corpus of articles divided into 5 sections: business, entertainment, politics, sports and technology. Different machine learning models such as naïve Bayes and support-vector machines were used as an attempt to make the supervised approach as effective as possible. However, the lack of an available local dataset meant that it was necessary to resort to a foreign one, instead. This approach produced poor results, as the significant difference in named entities caused certain documents to be misclassified. Hence, a different approach was chosen, namely the unsupervised topic-modelling technique called latent Dirichlet allocation (LDA). This was preferred due to its ability to detect and show latent topics using the term frequencies of the different articles, thus generating topics closely related to the dataset. It is also worth mentioning that a number of studies recorded the success of the LDA technique in modelling news topics [1, 3, 4].

Once the news articles were classified, the next step was to profile the authors of the comments. The two main challenges underlying this task included detecting topic relevance and analysing sentiment to understand how

Figure 1. Workflow diagram outlining the method of processing articles and comments for profiling news-portal users

users reacted to the article. With regard to the former, surface-based techniques based on cosine-similarity and Jaccard-similarity were adopted initially. However, results were not sufficiently productive due to the varying length of different articles and comments. Consequently, a different approach was identified, which was based on the LDA model created from the articles. This approach was applied to the comments, since it could recognise the most frequently used terms in order to categorise the topic accordingly [7]. To address the 'sentiment' aspect it was deemed best to utilise the gold-standard, lexicon-based approach of the sentiment-analysis tool Vader (Valence Aware Dictionary for Sentiment Reasoning) [2, 5]. This tool is specifically attuned to analyse comments posted on social media. It also proved successful in other domain contexts, including New York Times editorials, and hence deemed suitable for the purposes of this study.

To allow users to visualise the analysis of the article-comments relationship, a dashboard-style interface was developed, through which a researcher could select specific individuals who have submitted a comment, and receive analytical data about them. This data would include the user's overall sentiment, how topical the user's comments were and their favourite topic(s).

References/Bibliography:

[1] Anton Barua, Stephen W Thomas, and Ahmed E Hassan. What are developers talking about? an analysis of topics and trends in stack overflow. Empirical Software Engineering, 19(3):619–654, 2014.

[2] Clayton J Hutto and Eric Gilbert. Vader: A parsimonious rule-based model for sentiment analysis of social media text. In Eighth international AAAI conference on weblogs and social media, 2014.

[3] Carina Jacobi, Wouter Van Atteveldt, and Kasper Welbers. Quantitative analysis of large amounts of journalistic texts using topic modelling. Digital Journalism, 4(1):89–106, 2016.

[4] David Mimno, Hanna M Wallach, Edmund Talley, Miriam Leenders, and Andrew McCallum. Optimizing semantic coherence in topic models. In Proceedings of the conference on empirical methods in natural language processing, pages 262–272. Association for Computational Linguistics, 2011.

[5] Alejandro Moreo, M Romero, JL Castro, and Jose Manuel Zurita. Lexicon-based comments-oriented news sentiment analyzer system.Expert Systems with Applications, 39(10):9166–9180, 2012.

[6] Manos Tsagkias, Wouter Weerkamp, and Maarten De Rijke. Predicting the volume of comments on online news stories. In Proceedings of the 18th ACM conference on Information and knowledge management, pages 1765–1768, 2009.

[7] Jing Wang, Clement T Yu, Philip S Yu, Bing Liu, and Weiyi Meng. Diversionary comments under political blog posts. In Proceedings of the 21st ACM international conference on Information and knowledge management, pages 1789–1793. ACM, 2012.

Fast Approximation of Euclidean TSP

Nikola Glushkov **Supervisor: Dr Michel Camilleri**

Co-Supervisor: Dr. John Abela

Course: B.Sc. IT (Hons.) Computing and Business

Determining feasible routing paths is a substantial challenge in the humanitarian aid, civil, military, and commercial spheres. Specifically, finding a tour on a road network is a principal concern in planning logistics. The essence of this problem is characterized by the Travelling Salesman Problem (TSP). In a TSP, the salesman visits all cities exactly once. Joining the last city with the starting one produces a Hamiltonian cycle on the cities. One of the many applications of TSP is the Green Vehicle Routing Problem which considers fuel tank capacity and refueling stops [1].

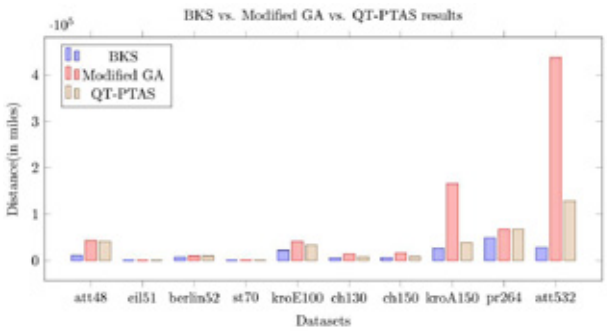

Figure 1. Best Known Solution (BKS) vs Genetic Algorithm (GA) vs Quad-Tree Polynomial-Time Approximation Scheme (QT-PTAS)

This study proposes a new effective approach that utilizes a Quad-Tree plane partitioning algorithm. Firstly, the Quad-Tree algorithm embeds all tour nodes along a plane. It then recursively partitions them into squares until each square contains at most three nodes. Secondly, the nodes within each square are repeatedly connected with each other in a bottom-up approach. The algorithm starts with the nodes in the smallest squares at the lowest level of the partitioning and works in a bottom-up approach until it reaches the largest child square of the Quad-Tree parent. This gives a Hamiltonian Cycle of the nodes (cities). This approach produces an approximation to the optimal tour length of the given nodes in polynomial time, thus becoming a quad-tree polynomial-time approximation scheme (PTAS). The final step of the study involves comparing this Quad-Tree PTAS with another sample solution based on a Genetic Algorithm (GA) approach.

The experiments are performed on ten real-life datasets of varying sizes. All the datasets are passed through the two approaches and the outcomes are analysed. The results obtained show that, in these experiments, the Quad-Tree PTAS consistently outperforms the GA in terms of tour length cost and computational time. The improvement in cost and time and the effectiveness of the Quad-Tree PTAS when compared to the GA approach grows with the size of the dataset. This may indicate that this improvement would increase with larger datasets. Therefore, the proposed Quad-Tree PTAS approach holds promise when compared with the Genetic Algorithm meta-heuristic and may be applied effectively to other problem domains.

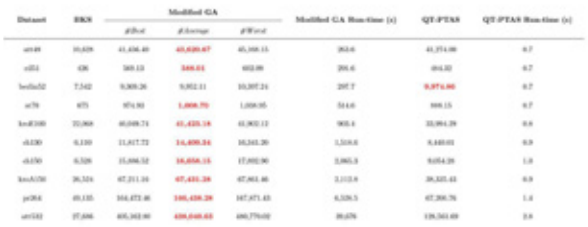

Figure 2. Genetic Algorithm (GA) vs Quad-Tree Polynomial-Time Approximation Scheme (QT-PTAS) results

References/Bibliography:

[1] S. Erdoğan and E. Miller-Hooks, "A Green Vehicle Routing Problem," Transportation research part E: logistics and transportation reciew, vol. 48, no.1, pp. 100-114, 2012.

Analysis of Aviation Safety and Aviation Accidents

Jamie Grech **Supervisor: Dr Joel Azzopardi**

Course: B.Sc. IT (Hons.) Artificial Intelligence

Aviation is a key method of transport in modern society and statistically also one of the safest [1]. Nevertheless, considering that there are tens of millions of commercial flights, annually [2], some accidents are inevitable.

Data mining is a process for sourcing new and potentially significant information from large databases. This makes it an invaluable tool in any field of research where data is recorded on a large scale, and the field of aviation safety is no exception. A multitude of data-mining techniques has already been employed for this field in previous research efforts, using a variety of databases. Some of the most notable examples are the aviation database of the National Transportation Safety Board (NTSB) of the United States, and the NASA Aviation Safety Reporting System (ASRS) database. Such databases could contain either structured data or unstructured data, such as reports written by airport staff – or a mix of the two. For the purpose of this project, the focus was on finding noteworthy correlations and deviations, primarily using structured data, as extracting information from the unstructured data would require the use of various natural language-processing techniques.

Passenger fatalities per billion passenger miles 2000–2009.

Riding a motorcycle	212.57
Driving or passenger in a car or light truck	7.28
Passenger on a local ferryboat	3.17
Passenger on commuter rail and Amtrak	0.43
Passenger on urban mass transit rail (2002–2009)[a]	0.24
Passenger on a bus (holding more than 10 passengers – transit, inter-city, school, charter)	0.11
Passenger on commercial aviation	0.07

Figure 1. Fatalities per billion passenger miles of different modes of transport in the US between 2000-2009 [1]

The research carried out in this project aims to contribute to the discovery of significant deviations in this data, thus supporting domain experts in increasing air-travel safety. This could be achieved by comparing the effectiveness of different contrast set mining algorithms on the data, based on a variety of metrics without relying on expert feedback. Moreover, the use of multiple databases would allow for a better comparison of the aforementioned algorithms, by allowing their evaluation on differently structured sets of data. Finally, this variety of data would also be of use in searching for a wider variety of different deviations, such as the differences between fatal and non-fatal accidents, or the differences between incidents caused by human factors and incidents occurring as a result of technical issues.

Rule	Length	Coverage	Group	Lift
('Clouds=>Yes', 'Flight Plan=>IFR', 'Cruise=>Yes', 'Initial Approach=>No')	4	14	Primary Problem=>Aircraft	3.703703704
('Clouds=>Yes', 'Light=>Daylight', 'Flight Plan=>IFR', 'Cruise=>Yes')	4	17.2	Primary Problem=>Aircraft	3.142857143
('Fog=>No', 'Light=>Dawn', 'Taxi=>No', 'Climb=>Yes')	4	17.7	Primary Problem=>Aircraft	3.037634409
('Flight Conditions=>IMC', 'Snow=>Yes', 'Taxi=>No', 'Climb=>Yes')	4	19.8	Primary Problem=>Aircraft	2.469623915
('Clouds=>Yes', 'Haze=>No', 'Flight Plan=>IFR', 'Cruise=>Yes')	4	19.2	Primary Problem=>Aircraft	2.463157895
('Icing=>No', 'Snow=>Yes', 'Turbulence=>Yes', 'Climb=>No')	4	12.3	Primary Problem=>Human Factors	2.309077278
('Flight Conditions=>VMC', 'Haze=>No', 'Light=>Dusk', 'Flight Plan=>IFR')	4	12.5	Primary Problem=>Aircraft	2.381395349
('Light=>Dusk', 'Mission=>Passenger', 'Taxi=>No', 'Climb=>Yes')	4	19.5	Primary Problem=>Aircraft	2.381226054
('Clouds=>Yes', 'Flight Plan=>IFR', 'Cruise=>Yes', 'Descent=>No')	4	26.9	Primary Problem=>Aircraft	2.37908497
('Icing=>No', 'Turbulence=>Yes', 'Parked=>Yes', 'Initial Climb=>No')	4	20	Primary Problem=>Human Factors	2.291666667
('Flight Conditions=>Marginal', 'Fog=>No', 'Flight Plan=>IFR', 'Cruise=>Yes')	4	25.6	Primary Problem=>Aircraft	2.021367521
('Icing=>Yes', 'Taxi=>No', 'Climb=>Yes')	3	25.8	Primary Problem=>Aircraft	2.019898504
('Light=>Daylight', 'Mission=>Ferry', 'Taxi=>No')	3	15.5	Primary Problem=>Aircraft	2.005218855

Figure 2. Sample contrast sets found from the ASRS database using the STUCCO algorithm

References/Bibliography:

[1] I. Savage, "Comparing the fatality risks in united states transportation across modes and over time, "Research in Transportation Economics, vol. 43, no. 1, pp. 9–22, 2013.

[2] https://www.statista.com/statistics/564769/airline-industry-number-of-flights (Last accessed 17/12/2019)

Remedi: A Medical Information Extraction System

Valerija Holomjova **Supervisor: Dr Charlie Abela**

Co-Supervisor: Mr Dylan Seychell

Course: B.Sc. IT (Hons.) Artificial Intelligence

Articles about modern advancements and discoveries in the medical industry are being published online on a daily basis. However, due to privacy issues, there is a lack of data available for building knowledge-extraction tools specialised in the medical industry. This research aims at creating a biomedical information-extraction tool called Remedi.

The system is specialised in identifying 'medical problem' and 'treatment' entities from unstructured text in the medical domain, as well as the relations between them. The tool consists of a biomedical named-entity recognition (BM-NER) model, as well as a biomedical relation-extraction (BM-RE) model. A subset of the i2b2 2010 challenge dataset [6] has been acquired to train the BM-NER and BM-RE components. The dataset includes annotated anonymised medical reports that could be used for concept extraction and relation-classification tasks.

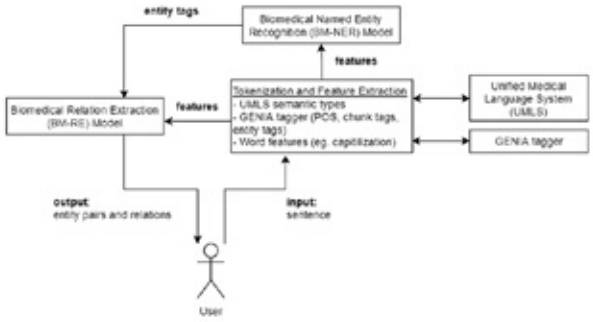

Figure 1. Dataflow diagram of the proposed system, Remedi

The research process revealed that most participants of the i2b2 concept-extraction challenge, which acquired the top results overall, employed conditional random field (CRF) models [6]. On the other hand, more recent research explores the use of neural networks, particularly, recurrent neural network (RNN) and convolutional neural network (CNN) models, which generated higher F-score results [7, 4]. Separate research explored the performance of a Bi-LSTM-CRF model on the dataset, which combined a bidirectional long-short-term-memory (Bi-LSTM) unit with a CRF layer for sequential tagging tasks, and obtained promising results [2]. However, Chalapathy et al. [2] did not mention the use of external sources (such as the Unified Medical Language System (UMLS) [1] and the GENIA tagger [5]) which has been noted to increase the performance of many models in the i2b2 challenge [6]. Hence, this research evaluates the impact of features extracted from external sources on the Bi-LSTM-CRF model, and compares it to the performance of a baseline CRF model using the same external features.

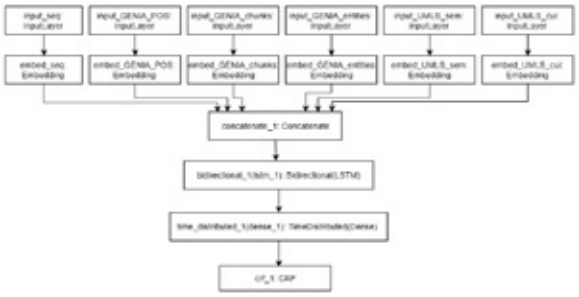

Figure 2. Structure of the Bi-LSTM-CRF model in Keras

For the i2b2 relation-extraction challenge, the highest results were obtained using support vector machines (SVM) [6]. Research by Chikka et al [3] explored the use of Bi-LSTM models in comparison to constructed SVM models using a subset of the i2b2 dataset, which was also applied in this study. However, Chikka et al. [3] also do not make reference to the use of external sources. As a result, this study explores the performance of a Bi-LSTM model for relation-extraction using the same external sources employed in the experimentation for the BM-NER component, and comparing it to the models constructed in the research. This dissertation also carries out hyperparameter tuning for all implemented models, to determine which set of parameters would yield the best results for each model.

References/Bibliography:

[1] Olivier Bodenreider. The unified medical language system (umls): integrating biomedical terminology. Nucleic acids research, 32(suppl1): D267–D270, 2004.

[2] Raghavendra Chalapathy, Ehsan Zare Borzeshi, and Massimo Piccardi. Bidirectional lstm-crf for clinical concept extraction. arXiv preprint arXiv:1611.08373, 2016.

[3] Veera Raghavendra Chikka and Kamalakar Karlapalem. A hybrid deep learning approach for medical relation extraction. arXiv preprint arXiv:1806.11189, 2018.

[4] Zengjian Liu, Ming Yang, Xiaolong Wang, Qingcai Chen, Buzhou Tang, Zhe Wang,and Hua Xu. Entity recognition from clinical texts via recurrent neural network. BMC medical informatics and decision making, 17(2):67, 2017.

[5] Y Tsuruoka. Genia tagger: Part-of-speech tagging, shallow parsing, and named entity recognition for biomedical text. Available at: www-tsujii. is. su-tokyo. ac.jp/GENIA/tagger, 2006.

[6] Özlem Uzuner, Brett R South, Shuying Shen, and Scott L DuVall. 2010 i2b2/va challenge on concepts, assertions, and relations in clinical text. Journal of the American Medical Informatics Association, 18(5):552–556, 2011.

[7] Yonghui Wu, Min Jiang, Jun Xu, Degui Zhi, and Hua Xu. Clinical named entity recognition using deep learning models. In AMIA Annual Symposium Proceedings, volume 2017, page 1812. American Medical Informatics Association, 2017.

Environmental Event Discovery through Time Series Anomalies

Andre Jenkins | **Supervisor: Dr Joel Azzopardi**

Co-Supervisor: Dr Adam Gauci

Course: B.Sc. IT (Hons.) Artificial Intelligence

Our planet has experienced numerous environmental events throughout its history, and there has been significant interest over the centuries to the present day to discover more about them when they happen.

Each event could be said to either be the result of abnormal meteorological behaviour, or cause further meteorological changes. A wildfire, for example, will always contribute to an increase in temperature, due to the fires involved. A flooding would most likely be caused by a substantial increase in rain that happened in the vicinity of the area.

Figure 1. Interpreting data to detect the 2019 wildfires in Australia

Most events have a primary cause or direct effect that could be detected, if sought or investigated. A drought is usually the result of a significant lack of rainfall over a long period of time, and so on.

The system being proposed in this study utilises time series data in order to detect anomalies, which are subsequently presumed to correlate to a specific event. It processes a dataset for the primary datapoint within it (temperature, precipitation, etc.), and then executes a discord-discovery algorithm over it in order to identify any anomalies. Once this is done, the discords could then be linked to a news article by utilising the time span of the anomaly, and some additional parameters, such as which query to submit. A group of articles is then requested using the application programming interface (API) of the relevant news source, the most relevant ones being those revolving around the event that occurred. For example, one could input year-long precipitation data for a certain city in the UK, enter 'UK flooding' as information request, select The Guardian as news source, to be presented with the available articles regarding a flooding in the city under review.

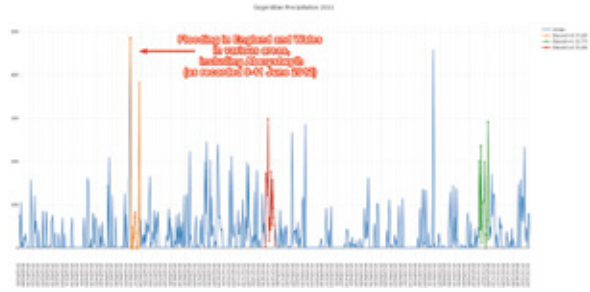

Figure 2. Interpreting data to detect the 2012 flooding in England and Wales

Although the current system is a prototype, it could easily be extended. One could implement APIs for further news sources, or attach the dataset to the system itself, in order that it could dynamically find the data and even generate the query itself, requiring even less input from the user. Observing the results obtained on temperatures linked to wildfires, and precipitation for instances of flooding, the system detected such events with very high accuracy, confirming that the significant anomalies in the time series data strongly suggest the occurrence of an associated event.

Automatic crime information gathering and data analytics from online news reports

Janica Spiteri | **Supervisor: Dr John Abela**

Co-Supervisor: Dr Peter Albert Xuereb

Course: B.Sc. IT (Hons.) Software Development

One of the major challenges encountered by law enforcement entities is the prioritisation and rostering of resources, with a view to maximise the chances of being equipped with the right resources at the right place, and at the right time.

This research proposes a hybrid machine learning technology that uses a set of customised web crawlers to gather data on a daily basis from newspaper articles. Articles that deal with criminal offences would be identified, analysed and their inherent details extracted using natural language processing (NLP) technology. Articles from different sources are converged using a standardised format that allows the details of the crime (such as nature of the crime, location, time, criminal, etc.) to be accessed easily. Related data, such as population, literacy etc. are also extracted from other sources using dedicated crawlers and cross-referenced with the events related to the actual crime. Web crawling is automated using a special bot designed to initiate the crawling processes regularly.

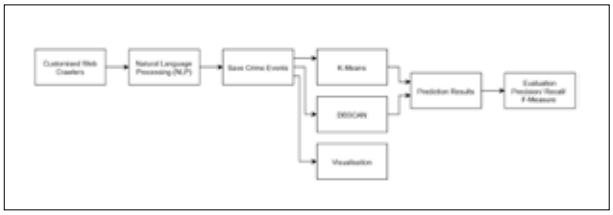

Figure 1. Block diagram of the system

This study proposes a visualisation engine that would allow users to quickly and effectively browse the criminal-event database using a feature-rich search engine that would enable specific parameters to be easily identified and depicted. Various representations of the resulting data (such as geographical heat maps, graphs, calendar heat maps, etc.) are possible,

Previous research in similar areas has utilised various machine learning techniques with varying success rates. This research aims to study the effectiveness of k-means and DBSCAN-based technologies [1] when applied to crime prediction. While k-means uses a purely statistical past-data based model to attempt to predict the incidence of crime, DBSCAN uses clustering techniques, which could include other datasets in addition to past criminal event data.

Various data has been used to evaluate the performance of the proposed technology – with encouraging results. The precision / recall / F-measure technique used in previous studies [2], [3], has been utilised to compute the accuracy of both techniques. Moreover, geographically distant areas (Malta and Boston), where used to evaluate different crime patterns. Although the large number of possible prediction configurations make it very difficult to cover all the possible scenarios, both techniques performed quite well. In particular, the k-means technique proved to be slightly more accurate in predicting recurring crimes.

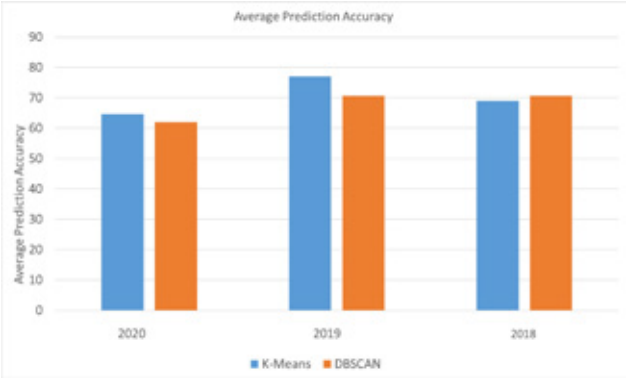

Figure 2. Average prediction accuracy chart

References/Bibliography:

[1] Baboo, S. Santhosh. "An enhanced algorithm to predict a future crime using data mining." Int J Comput Appl 975: 8887 (2011).

[2] Iqbal, Rizwan, et al. "An experimental study of classification algorithms for crime prediction." Indian Journal of Science and Technology 6.3: 4219-4225 (2013).

[3] Almanie, Tahani, Rsha Mirza, and Elizabeth Lor. "Crime prediction based on crime types and using spatial and temporal criminal hotspots." arXiv preprint arXiv:1508.02050 (2015).

Movie recommendations using machine learning algorithms

Steve Spiteri | **Supervisor: Dr Clyde Meli**

Co-Supervisor: Dr John Abela

Course: B.Sc. IT (Hons.) Software Development

Data Science

In today's advancing technology, the use of the internet is growing day by day. Hence, has become more important than ever that internet users acquire the tools to help them process the huge amounts of data on the Web. This brings about the need for recommender systems to suggest items that would be of possible interest to the individual users. This research project focused on the task of recommending films to users through machine learning algorithms, rather than methods like Collaborative Filtering And Content Based Filtering [1].

The two most popular types of recommender systems are content-based and collaborative filtering. These methods are implemented in different scenarios of recommendation, where each method is implemented differently to be compatible with the corresponding system. It is to be noted that these methods also have their drawbacks. Content-based filtering operates by recommending an item that is similar to an item that the user has already rated. If the item would not contain enough information, or if the user would not have rated any items, the recommendation would be inaccurate [1]. Moreover, collaborative filtering (CF) is based on the user's historical preference on a set of items. It uses rating information from all other users to provide predictions for user-item interactions [2]. One of the disadvantage of this method is that it suffers from the cold-start problem, which is a start-up issue that could occur in this context when a new user and/or item would be entered without any ratings, and the recommender system would not have enough data to make any reasonable inferences [3].

The above-mentioned difficulty provided the main motivation for utilising machine learning (ML) algorithms to recommend films for the purpose of this project. ML algorithms are applicable to all kinds of data and could be employed without the need to custom-build the algorithm for the particular scenario. This study is based on determining whether using various ML algorithms would lead to better rating results when compared to

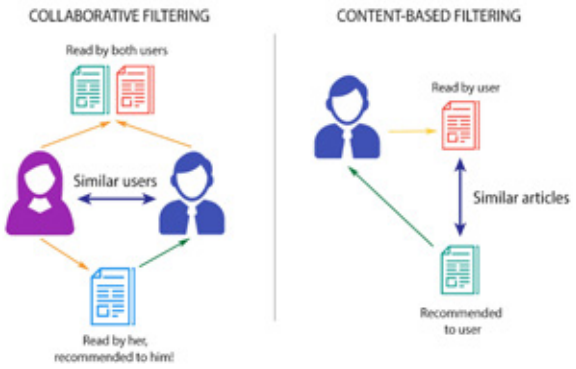

Figure 1. Different recommendation-system approaches [4]

the use of collaborative or content-based filtering. ML models could achieve better results because the model could be trained by using more attributes (features) about the item (i.e. a film, in this case) and the user. For example, in the case of users, one could add age, location and occupation, whereas for the item one could add the actors appearing in the film or the film genre. This information enhances the recommender system because the ML algorithm is also considering these attributes, whereas other algorithms, such as collaborative filtering, are not equipped to include them.

The results of this research were evaluated by using mean absolute error (MAE), by measuring the errors between paired observations and comparing them to other researches. The research led that, when training a model using the whole dataset, the results were not as quite as good as to other research. Moreover, the ML model was implemented to focus on each individual film rather than all the films at once, and an ML model was trained on the dataset of each film, thus generating better results. This outcome led to the conclusion that an ML model offers better recommendations when trained on individual items.

References/Bibliography:

[1] Anonymous "Cold start (recommender systems)," 2020. Available: https://en.wikipedia.org/w/index.php?title=Cold_start_(recommender_systems)&oldid=973940047.

[2] B. Rocca. (-06-12T07:48:22.390Z). Introduction to recommender systems. Available: https://towardsdatascience.com/introduction-to-recommender-systems-6c66cf15ada.

[3] N. Sharma. (-02-05T16:02:43.917Z). Recommender Systems with Python — Part I: Content-Based Filtering. Available: https://heartbeat.fritz.ai/recommender-systems-with-python-part-i-content-based- filtering-5df4940bd831.

[4] S. Doshi. (-02-20T16:32:35.531Z). Brief on Recommender Systems.

Using unsupervised and supervised machine learning to discover discrepancies between the two counter-circulating beams of the Large Hadron Collider

Michael Vella Zarb **Supervisor: Dr Ing. Gianluca Valentino**

Course: B.Sc. (Hons.) Computing Science

The function of the Large Hadron Collider (LHC) is to allow hadrons (e.g. protons, neutrons) to collide at great speed, for the purpose of analysing the outcomes in order to gain a greater understanding of matter at the subatomic level. To achieve this, the LHC has been designed to consist of two circular, concentric beam pipes, each containing a particle beam circulating in opposing directions. There are four points at which the beam pipes cross over to bring about collisions. Superconducting magnets are used to direct the beams into their circular path along the pipes, and to focus the beams before the crossover points to increase the likelihood of collisions.[1]

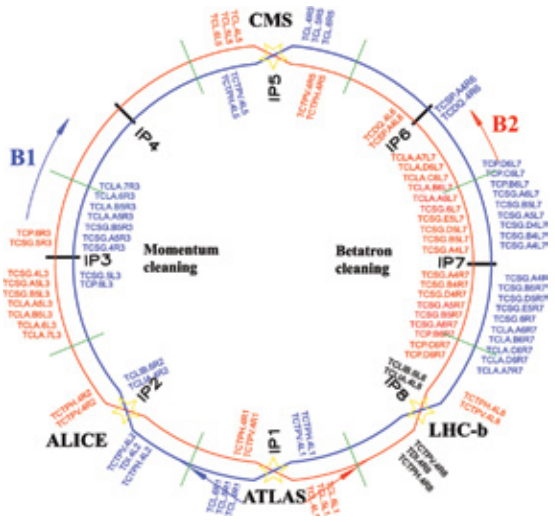

Figure 1. Collimator layout of the LHC with beam-loss monitors placed at the collimators. (source Redaelli, Stefano 2016)

A cross-section of the beam pipe reveals that the particles form a normal distribution, with most of the particles being drawn towards the centre of the beam pipe. The particles furthest away from the centre are collectively known as the beam halo. Special devices - collimators [2] are placed at certain points around the beam pipes to absorb the beam halo. If not absorbed, the beam halo could damage equipment, e.g. superconducting magnets might be quenched, or could cause noise obfuscating the desired results concerning the beam collisions.[3] Beam-loss monitors (BLMs) are therefore used to detect the presence of a beam halo. If dangerously high losses would be detected, the beams could be terminated.[4]

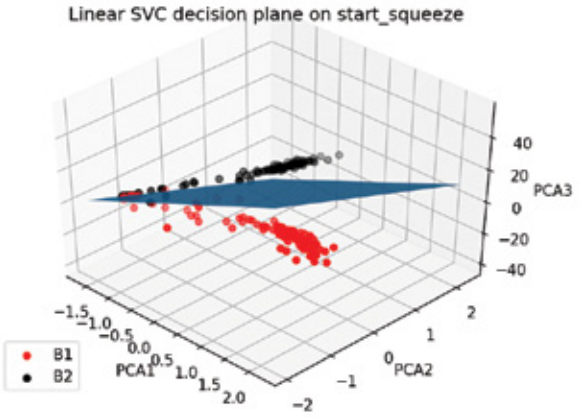

Figure 2. 3D graph illustrating how linear SVM could identify / delineate a decision boundary, which separates the Beam 1 losses from the Beam 2 losses almost perfectly

In theory, LHC beams should be symmetric, but this is not necessarily the case in practice. Symmetric BLMs (e.g.TCP.6R3 and TCP.6L3) were paired into single features, and supervised and unsupervised machine learning techniques were used (KMeans, DBSCAN clustering, linear SVM) to discover discrepancies between the BLM readings of the two beams. Each technique varied in its ability to separate the BLM data into two clusters (or classes on either side of a decision boundary, in the case of linear SVM.) A low separability would suggest minimal discrepancies between the losses of the two beams. The technique that could separate the data most efficiently was linear SVM, almost perfectly separating (and therefore classifying) the data into Beam 1 losses on one side of the decision boundary, and Beam 2 losses on the other. This indicates that discrepancies between the two beams are indeed significant.

References/Bibliography:

[1] https://lhc-machine-outreach.web.cern.ch/lhc-machine-outreach/beam.htm
[2] https://home.cern/news/news/experiments/collimators-lhcs-bodyguards
[3] https://home.cern/news/news/accelerators/crystal-cleaning-lhc-beam
[4] https://www.worldscientific.com/doi/pdf/10.1142/S0217751X13300354 4.11.2.

Computational Mobility in JavaScript

Neil Borg **Supervisor: Dr Kevin Vella**

Course: B.Sc. (Hons.) Computing Science

The aim of this project is to build and evaluate a JavaScript framework that would facilitate the migration of both code and execution state in a distributed computational environment. The framework targets multiple nodeJS instances, which can operate across both multicore and distributed systems. The said framework seeks to streamline the implementation of distributing an existing system, by allowing the dynamic migration of JavaScript functions across nodeJS instances as they execute. The framework also features a remote-execution facility, allowing the migration of both code and execution state across nodeJS instances. This facility paved the way for the construction of an in-built, load-balancing mechanism.

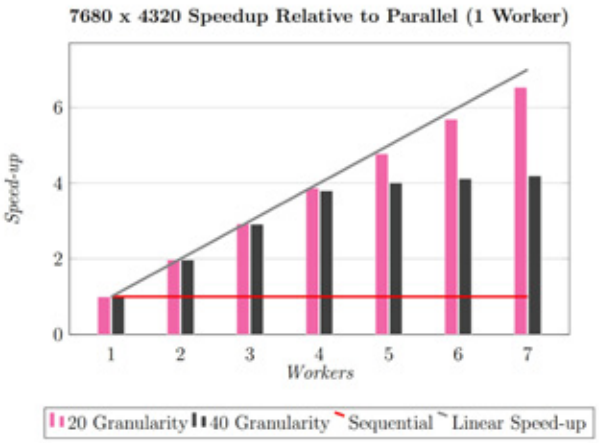

Figure 1. Speed-up of the 7680 x 4320 Mandelbrot fractal

While JavaScript executing within a single nodeJS instance only harnesses a singular core [1], the ability of the framework to shift computation between nodeJS instances in a distributed multiprocessor environment provides an opportunity to exploit the inherent parallelism that is available. In order to demonstrate the effectiveness of the framework in harnessing resources for parallel computation, two task-farming case studies were written using the framework. The first case study involved a Mandelbrot set solver [3] that performs iterative calculations on the complex plane in order to generate an image representation, and which represents a highly unbalanced computational workload. The second case study involved a synthetic benchmark representing a highly balanced computational workload, for comparative purposes.

Both of these scenarios were deployed on an octa-core computer, which was set to run a nodeJS instance on every core. Results were taken over a broad range covering both the problem size and task granularity. Speed-up and efficiency-performance metrics were derived to provide concrete evidence that the framework functions as intended.

An analysis of the results concluded that the observed speed-up was bound by the problem size: small images were not worth splitting up for distributed computation. On the other hand, large images did not encounter this limit, and therefore resulted in higher degrees of speed-up. The outcomes are comparable to Gustafson's predictions for parallel computing [2]. Finer-task granularity was observed to aid the speed-up and efficiency metrics by exploiting the load-balancing characteristics of the system.

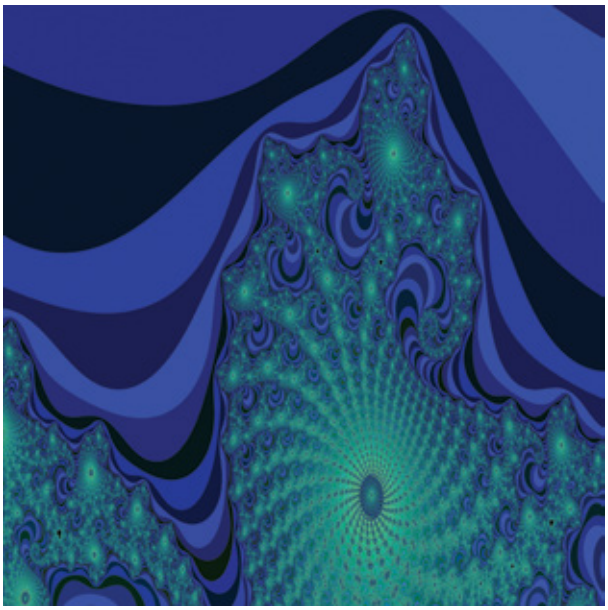

Figure 2. Result from computation of the Mandelbrot fractal

References/Bibliography:

[1] K. Ghadyani. The reality of JavaScript performance, Mar 2019.

[2] J. L. Gustafson. Gustafson's Law, pages 819–825. Springer US, Boston, MA, 2011.

[3] B. Mandelbrot, P. Jones, C. Evertsz, and M. Gutzwiller. Fractals and Chaos: The Mandelbrot Set and Beyond. Selecta (Springer). Springer, 2004.

Memory Forensics of Qakbot

Steve Borg **Supervisor: Dr Mark Joseph Vella**

Course: B.Sc. IT (Hons.) Software Development

In their ongoing pursuit of developing malware, a common technique used by malware authors is process injection, through which malicious code is injected into benign processes with escalated privileges. In the past, signature-based detection was generally considered to be sufficient for malware detection. However, with the advent of polymorphism [1] – which became one of the most prevalent detection-evasion techniques – anti-virus signatures were no longer effective, due to the malware's ability to change its appearance virtually at will.

Qakbot [2] is a prime example of malware that, notwithstanding the numerous signatures [3,4] written over the years, continued to evolve and evade detection [5]. Consequently, the detection of the Qakbot sample would be delayed [6], necessitating the use of digital investigation tools for incident response. This malware has evolved and managed to blend into regular Windows processes (see Figure 1), emphasising the importance of memory forensics in identifying the exact workings of Qakbot [2] and the necessity of reconstructing the timeline of events occurring after the malware infection. A prominent obstacle to the analysis of the Qakbot malware is that it includes a packing layer [7] that compresses parts of the malware to avoid detection, and thus hindering analysis.

In this study, reverse software engineering (RSE) [8] and dynamic binary instrumentation (DBI) [9] techniques were used to produce forensic tools to assist incident responders in identifying precisely the processes being created and potentially injected. The first tool that was developed was based on state-of-the-art system logs and memory forensics. This would provide a very reliable indication, but not watertight proof, that the benign process being injected is indeed mobsync.exe. The second tool is very similar to the first, with the addition of volatility filters that attempt to detect the newly discovered information from advanced analysis. The third and final tool that was developed is a custom tool based on DBI

Figure 1. Qakbot Kill Chain

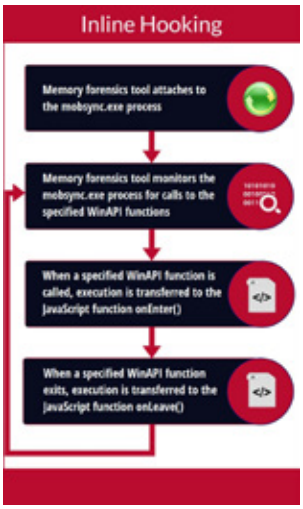

Figure 2. Dynamic Binary Instrumentation Through Inline Hooking

(see Figure 2) which, through partial but timely memory dumps, could obtain the necessary infection evidence. The complete mobsync.exe misuse picture comes at the expense of computer memory and storage overheads.

References/Bibliography:

[1] Lord, N. (2018, September 11). What is Polymorphic Malware? A Definition and Best Practices for Defending Against Polymorphic Malware. Retrieved from https://digitalguardian.com/ blog/what-polymorphic-malware-definition-and-best-practices- defending-against-polymorphic-malware

[2] Falliere, N. (2009). W32.Qakbot in Detail. Retrieved from https://web.archive.org/web/20191010084219/http://www. symantec.com/content/en/us/enterprise/media/security_ response/whitepapers/w32_qakbot_in_detail.pdf

[3] Barabosch, T. (2017, April 21). qakbot.yara. Retrieved from https://github.com/tbarabosch/quincy-complementary-material/ blob/master/yara/qakbot.yara

[4] Yara Rules. (2019, July 5). Retrieved from https://malpedia.caad.fkie.fraunhofer.de/details/win.qakbot

[5] Benge, A. & Randolph, N. (2019, May 2). Qakbot levels up with new obfuscation techniques. Retrieved from https://blog.talosintelligence.com/2019/05/qakbot-levels-up-with-new- obfuscation.html

[6] 2018-12-14 - Emotet infection with Qakbot. (2018, December 14). Retrieved from https://www.malware-traffic-analysis.net/2018/12/14/index.html

[7] K2 Intelligence. (2017, April 28). Qakbot: Insight into a Persistent, Polymorphic Malware from the Cyber Defense Team at K2 Intelligence. https://www.jdsupra.com/legalnews/qakbot- insight-into-a-persistent-26871/

[8] Introduction to Software Engineering/Reengineering/Reverse Engineering. (n.d.). Retrieved from https://en.wikibooks.org/wiki/Introduction_to_Software_ Engineering/Reengineering/Reverse_Engineering

[9] Dynamic Binary Instrumentation. (2007). Retrieved from http://uninformed.org/index.cgi?v=7&a=1&p=3

The Textures Shading Cache - An Approach to Asynchronous Interactive Illumination

Aidan Cauchi **Supervisor: Dr Sandro Spina**

Co-Supervisor: Dr Keith Bugeja

Course: B.Sc. (Hons.) Computing Science

In computer graphics rendering refers to the process of synthesizing an image from a mathematical description of a scene. Traditionally, interactive applications based on GPU-accelerated rasterization utilize techniques that approximate the effects of light within a scene, sacrificing correctness (realism) in favour of maintaining interactivity [2, 3, 4]. In order to increase the realism of rendered images, light-transport algorithms that correctly simulate the propagation of light in a scene are increasingly being used in interactive rendering. These algorithms push the computational bottleneck in interactive visualization towards the shading of pixels in an image and can, especially on low-end GPUs, negatively impact interactivity.

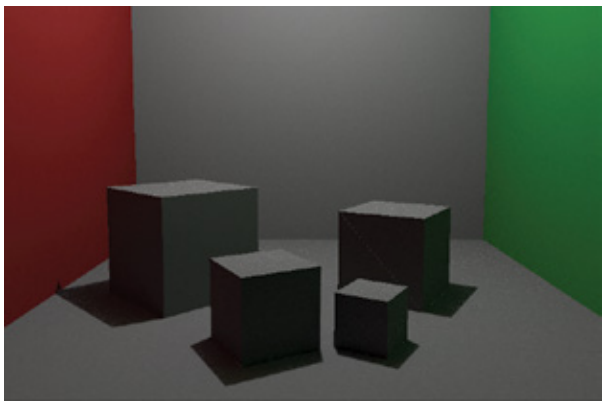

Figure 1. Basic scene, rendered using the method followed [or] adopted in the study

This project explores the effects of decoupling physically correct shading computations from the traditional rasterization-based rendering process, through the creation of a texture cache (see Figure 2). This cache is kept updated asynchronously at run-time. A simple scheduler determines which shading computations are to be computed asynchronously, and which would be executed in lock-step with frame rendering, through the traditional approach. This was done with the aim of achieving a stable frame rate without dips (sudden drops in frame rate) or stuttering (irregular delays between frames), whilst preserving image quality. A solution was implemented in C++ using the OpenGL API for the traditional renderer. Moreover, a path tracer [1] was built using Nvidia Optix for the asynchronous calculations. A texture atlas is used as the texture cache, where each primitive in the scene is allocated a space within this atlas during run-time.

The proposed solution was evaluated thoroughly to assess the effectiveness of the concept with respect to both performance and image quality. Three modes were evaluated, in particular: decoupling of all the shader computations as illustrated in Figure 1; decoupling only of the texture and bounce-lighting computations; decoupling only the bounce-lighting calculations. The performance of each mode was evaluated, measuring both frames per second and memory usage. Furthermore, the quality of the images produced was also determined using standard image-quality metrics HDR-VDP and SSIM against high-quality versions of the scenes rendered offline.

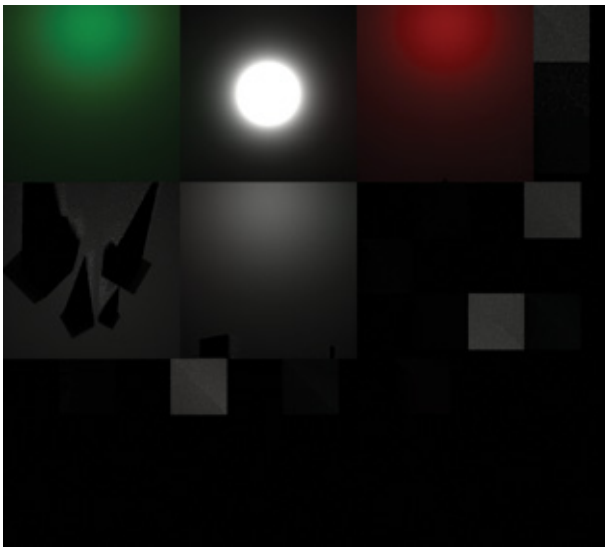

Figure 2. Texture cache output for the scene portrayed in the accompanying image

References/Bibliography:

[1] James T. Kajiya. 1986. The rendering equation. SIGGRAPH Comput. Graph. 20, 4 (Aug. 1986), 143–150. DOI:https://doi.org/10.1145/15886.15902

[2] Bui Tuong Phong. 1975. Illumination for computer generated pictures. Commun. ACM 18, 6 (June 1975), 311–317. DOI:https://doi.org/10.1145/360825.360839

[3] Lance Williams. 1978. Casting curved shadows on curved surfaces. SIGGRAPH Comput. Graph. 12, 3 (August 1978), 270–274. DOI:https://doi.org/10.1145/965139.807402

[4] Tobias Ritschel, Thorsten Grosch, and Hans-Peter Seidel. 2009. Approximating dynamic global illumination in image space. In Proceedings of the 2009 symposium on Interactive 3D graphics and games (I3D '09). Association for Computing Machinery, New York, NY, USA, 75– 82. DOI:https://doi.org/10.1145/1507149.1507161

A Mobile Application Providing a Mediterranean Diet Based on Physical Attributes and Weight Goal

Luka Djincharadze **Supervisor: Dr Lalit Garg**

Course: B.Sc. IT (Hons.) Computing and Business

Dietary factors could play a harmful or protective role in the case of chronic conditions [1], cancer risk [2] and even mental health [3]. However, research suggests that exploring generally healthful dietary information could be a complicated process, given the diverse array of evolving and widely available information, as well as aspects of diet that vary from one person to another. [4], In this regard, mobile nutrition applications have the potential to convey dietary knowledge in a convenient and scalable manner. [5] Yet, according to S. Scerri et al. [6], they could be pedantic and counter-productive if a burden is placed on the user.

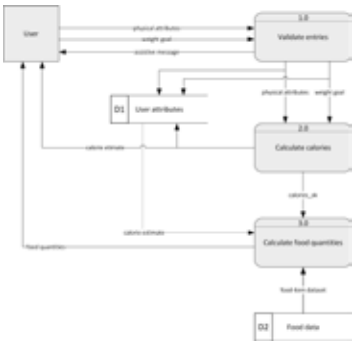

Figure 1. Level 1 DFD of the application

This study examines dietary research, mobile nutrition and application-design concepts for developing and evaluating a mobile diet application with simplicity, and the facilitation of making healthy food choices in mind. The overarching methodology was influenced by the waterfall software development model. This entails surveying before designing the application, subsequently undertaking user-acceptance testing (UAT) to test and evaluate the application after its implementation. Both the survey and UAT targeted University of Malta staff and students, although the survey also targeted Junior College staff.

The application itself seeks to provide a healthy Mediterranean diet, using an underlying deterministic model.

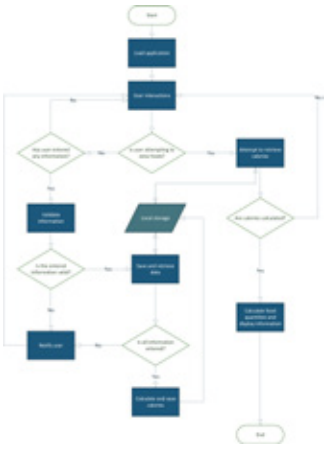

Figure 2. Flowchart of the application

The developed deterministic model behind the application utilises several user attributes and a weight goal to obtain a caloric estimate, towards subsequently calculating food quantities of eight different food groups according to the user's estimated calories and the compiled Mediterranean-diet food distribution and food dataset. The design of the application was influenced by Material Design for Android, with particular attention to simplicity, user experience, and the promotion of healthy food choices. The application is evaluated primarily according to these criteria. While aiming to facilitate the choice of healthy foods, the application does not seek to be tailor-made for users from a medical point of view and does not capture medical information.

According to the initial evaluation, the application was expected to be well-received, and the UAT findings did indeed present the application in a positive light (SUS score of 87.5 ± 8.75 at 99.9% confidence interval places the application in the upper quartile [7]). The majority of UAT participants preferred the developed application over others that expect constant and precise food inputs, and each of the nine participants responded in the affirmative when asked about the likelihood of using it.

References/Bibliography:

[1] A. Afshin et al., "Health effects of dietary risks in 195 countries, 1990–2017: a systematic analysis for the Global Burden of Disease Study 2017," The Lancet, vol. 393, (10184), pp. 1958-1972, 2019. Available: https://www.sciencedirect.com/science/article/pii/S0140673619300418. DOI: 10.1016/S0140- 6736(19)30041-8

[2] WHO, "Diet, nutrition and the prevention of chronic diseases," World Health Organization Technical Report Series, vol. 916, pp. i, 2003. Available: https://www.ncbi.nlm.nih.gov/pubmed/12768890.

[3] E. Frazão, "High costs of poor eating patterns in the united states," in America's Eating Habits: Changes and Consequences, USDA Economic Research Service, 1999, pp. 5-32.

[4] C. M. Pollard et al., "Who Uses the Internet as a Source of Nutrition and Dietary Information? An Australian Population Perspective," Journal of Medical Internet Research, vol. 17, (8), pp. e209, 2015. Available: https://www.ncbi.nlm.nih.gov/pubmed/26310192. DOI: 10.2196/jmir.4548.

[5] J. Rivera et al, "Mobile Apps for Weight Management: A Scoping Review," JMIR mHealth and uHealth, vol. 4, (3), pp. e87, 2016. Available: https://www.ncbi.nlm.nih.gov/pubmed/27460502. DOI: 10.2196/mhealth.5115.

[6] S. Scerri et al., "Understanding human-device interaction patterns within the context of mobile nutrition," in Dec 2015, pp. 1-7.

[7] A. Bangor, P. T. Kortum and J. T. Miller, "An Empirical Evaluation of the System Usability Scale," International Journal of Human-Computer Interaction, vol. 24, (6), pp. 574-594, 2008. Available:

Synthesis and Editing of Realistic Urban Road Networks

Gabriel Thomas Gauci | **Supervisor: Dr Sandro Spina**

Course: B.Sc. (Hons.) Computing Science

The 3D modelling of realistic road networks is beneficial to a number of areas of activity, including urban planning and traffic management [1][2]. Simulations carried out within these environments are key to the development and testing of software intended for driving autonomous vehicles [3][4][5]. The generation of these 3D models is a laborious task, entailing considerable time spent in ensuring that the environments created would be an accurate representation of the real world. The aim of this project was to design and develop a system to bootstrap this task by procedurally generating road geometry that fuses information from multiple, freely available sources.

Blender is a popular open-source 3D modelling software [6][7], with extensive support for scripting [8]; all functionality available through its graphical user interface (GUI) is exposed through its scripting application programming interface (API). This positions Blender as an ideal candidate for the procedural generation of geometry, which could be further refined within the GUI. Moreover, since Blender scripting is carried out in Python, a number of libraries were readily accessible for carrying out other important tasks for this project, such as network data queries/downloads and image processing. The main source of mapping information for this project was OpenStreetMap (OSM), which is built by a community of mappers who contribute data about roads and buildings from around the world [9]. Since OSM does not provide satellite views of the urban environments under consideration, satellite images from Google Maps [10] were used and automatically aligned to the procedurally generated road geometry in Blender. Image processing was carried out on these colour images to calculate the actual width of the various roads and the approximate location of any greenery present [11][12]. Since elevation data could not be inferred from a two-dimensional view, Google's Elevation API was used to adjust road geometry according to the retrieved data [13].

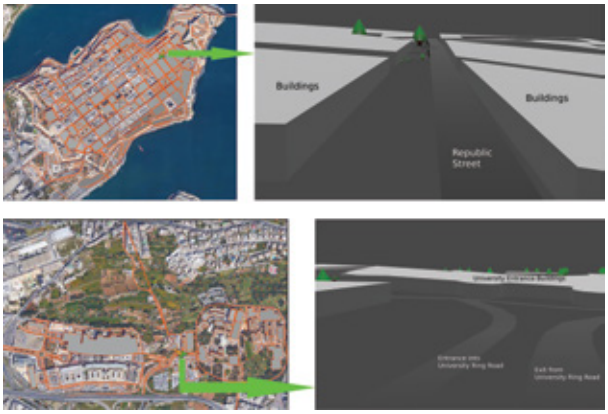

Figure 1. Top-down and street views of procedurally generated geometry for Valletta(top) and the University of Malta(bottom); generated geometry is rendered in Blender (right) and shown in the top-down view overlaid on the satellite image (left).

Upon its implementation, the evaluation of the proposed solution was found to deliver very accurate results which, in some cases, improved upon pre-existing solutions such as BlenderGIS [14]. Five locations covering a variety of urban layouts – namely: Ħ'Attard (San Anton area); Marsa (Royal Malta Golf Club area); Naxxar Centre, the University Ring Road and Mater Dei Hospital, and Valletta – were chosen to determine the quality of the generated geometry. Figures 1 and 2 illustrate the procedurally generated geometry for Valletta, and the University Ring Road and Mater Dei Hospital respectively (top view/street view in all cases). Future work includes the development of methods to procedurally add buildings, and to apply materials in order to increase the realism of the generated results.

References/Bibliography:

[1] R. Hamaina, T. Leduc, and G. Moreau. A structural analysis of the streets' network to urban fabric characterization. volume 2012, 07 2011.

[2] F. Molinari, A. Dethof, and J. Raisch. Traffic automation in urban road networks using consensus-based auction algorithms for road intersections. pages 1-8, 12 2018.

[3] M. DeBord. A waymo engineer told us why a virtual-world simulation is crucial to the future of self-driving cars, 2018.

[4] D. Silver. Simulation becomes increasingly important for self-driving cars, 2018.

[5] A. Dosovitskiy, G. Ros, F. Codevilla, A. Lopez, and V. Koltun. CARLA: An open urban driving simulator. pages 1-16, 11 2017.

[6] Blender. https://www.blender.org/. Accessed: 2020-04-24.

[7] Blender License. https://www.blender.org/about/license/. Accessed: 2020-04-24.

[8] Blender Features. https://www.blender.org/features/. Accessed: 2020-04-24.

[9] OpenStreetMap. https://www.openstreetmap.org/. Accessed: 2020-04-24.

[10] Google Maps. https://www.google.com/maps. Accessed: 2020-04-24.

[11] Changing Colorspaces. https://docs.opencv.org/3.4.8/df/d9d/tutorial_py_ colorspaces.html. Accessed: 2020-04-24.

[12] Color Space Conversions. https://docs.opencv.org/master/d8/d01/group__imgproc__color__conversions.html. Accessed: 2020-04-24.

[13] Google Maps Platform - Elevation API. https://developers.google.com/maps/ documentation/elevation/start. Accessed: 2020-04-21.

[14] domlysz. BlenderGIS. https://github.com/domlysz/BlenderGIS. Accessed: 2020- 04-21.

Voting Machine for People with Disabilities

Martina Gauci | **Supervisor: Dr Colin Layfield**

Co-Supervisor: Dr Conrad Attard

Course: B.Sc. IT (Hons.) Software Development

At present, the methods of casting a vote in a Maltese election is by paper ballot, Braille ballot or an audio guide. If a person is unable to vote through one of these options, the assistance of an electoral representative is offered, thus compelling the citizen to reveal the intended vote to the assistant. According to the Universal Declaration of Human Rights [1], all citizens are entitled the right of a secret vote. Therefore, it is of the utmost importance to provide a method that would offer independent voting for these citizens.

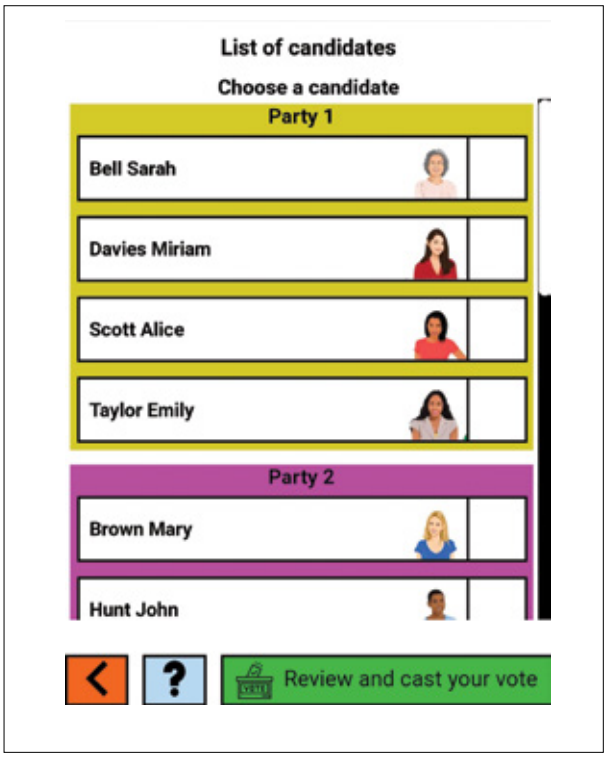

Figure 1. Ballot design for the final prototype

There are two main aspects to creating a voting machine: interface design and security. Due to time constraints, this project will only target the former. The aim of this study was to design an interface for a voting machine that would allow independent voting by those who are not able to do so using any of the methods currently available.

The artefact has been developed in an iterative and participatory design process. Initially, three wireframe designs were created. With the involvement of experts in accessibility, the low-fidelity prototypes were refined into more accessible designs and developed into web applications. A user- centric approach was taken to further improve the prototypes and users from the target audience evaluated the mid-fidelity and high-fidelity prototypes. The measurements extracted from established through the testing sessions involving the users were: effectiveness, efficiency and user satisfaction. These criteria enabled the identification of any significant differences between the usability of the prototypes. Feedback provided by the participants was used to support these findings. Since the testing was carried out by persons with disability, the usability measurements would reflect the accessibility of the artefacts.

Using the high-fidelity prototype, participants were able to cast the intended vote in under five minutes. All testers were also able to make the instructed changes to the vote. All prototypes obtained a good user satisfaction score, the highest mean score being 95.4 for the final prototype. The high-fidelity prototype, which was tested on a small number of participants, showed an improvement in effectiveness, efficiency and satisfaction although no statistical significance can be claimed.

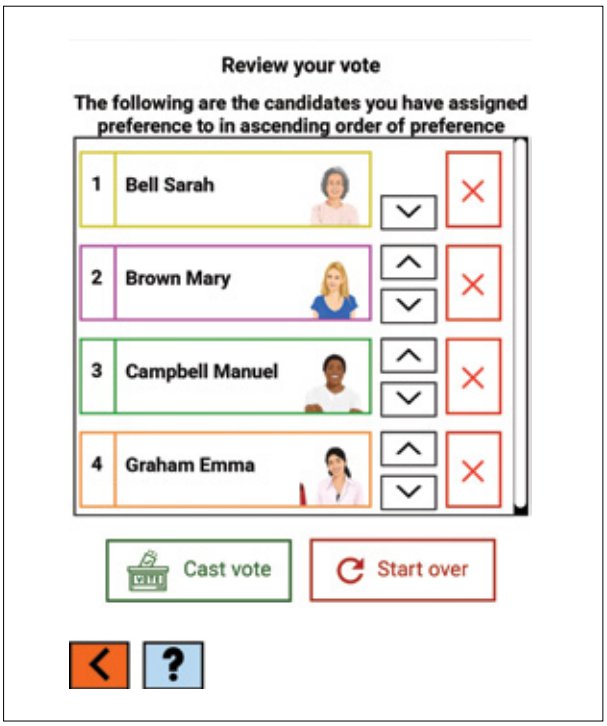

Figure 2. Review page for the final prototype

References/Bibliography:

[1] United Nations General Assembly, "Universal Declaration of Human Rights," 1948.

Teaching children with disabilities core literacy skills more effectively using a tablet-based context-driven web application

Juan Mizzi **Supervisor: Dr Peter Albert Xuereb**

Course: B.Sc. IT (Hons.) Software Development

Education is one of the pillars of any society – not least in Malta – since it brings progress and growth [1]. Over recent years, the Maltese education authorities have been implementing an emergent curriculum in the classrooms across the country. Emergent curriculum is a way of planning curriculum that is based on the children's interest and passion at a certain point in time [2]. To achieve this, it would be necessary for the educator to listen to and observe the child, to be able to plan around the child's interest, 'scaffold' building upon existing knowledge. Apart from this, the educator would also be expected to document the child's progress and communicate the child's progress with the parents to give them the opportunity of being involved in their child's education. While this process is intended towards improving the educational system, it is nevertheless time-consuming and demanding. Therefore, it might prove to be challenging to the educators to achieve the desired outcome with little or no assistance.

Taking the above into account, this project creates a platform with the aim of facilitating the educators' day-to-day tasks in the classroom. This is done by applying gamification to motivate the children to learn the letters of the alphabet in an engaging way that would be in tune with their interests, while simultaneously documenting and reporting the children's achievements to parents and the educators automatically. In this way, educators would have more time to observe, guide and assist the children during their learning process. Furthermore, the proposed platform would give parents and educators a clear view of the child's progress, making it easier for them to plan subsequent exercises to continue building on the child's knowledge.

Building upon the aid of professional educators, the application was developed using ASP.NET Core MVC, and was subsequently tested by occupational therapists and parents who assist children with learning disabilities. The methodology used was iterative, so that the website could be developed and improved through the feedback

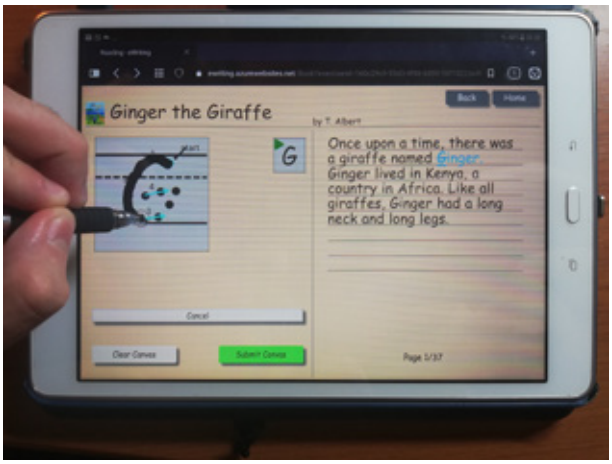

Figure 1. Using the tablet-based web application to teach children core literacy skills

obtained, and decisions agreed on previously. The participating parents felt that this platform motivated them to be more involved in their child's education since the platform is easy to use and easily accessible through electronic devices.

The occupational therapists (OTs) taking part in the project confirmed that this platform not only facilitated their work, but also allowed them more time to observe and guide children under their care. By recording and visually displaying the children's progress, it enabled them to obtain a good overview of what they know and the progress they are making, thus facilitating their forward-thinking process. The OTs also found that, by providing the appropriate context for the writing of letters of the alphabet – and because the children preferred using technological devices to pen and paper – the platform helped them to teach their pupils in a more engaging way. Furthermore, the OTs expressed the view that the platform enables a good handover, where detailed information about the child is efficiently passed on to other staff members from one year to the next. This better serves the child's interests.

References/Bibliography:

[1] O. Bonnici, "Education: a pillar of our new economic vision," 14 August 2020. [Online]. Available: https://www.independent.com.mt/articles/2020-08-14/blogs-opinions/Education-a- pillar-of-our-new-economic-vision-6736226107.

[2] E.K Ludman, "The Emergent Curriculum," Journal of Nutrition for the elderly, 1983.

Automated requirements engineering systems

Matthew Schembri | **Supervisor: Prof. Ernest Cachia**

Course: B.Sc. IT (Hons.) Software Development

Requirements elicitation and analysis is a development phase carried out at the initial stages of the development life cycle of every professional software solution. Historically, and by virtue of its nature, this development phase is prone to subjectivity and lack of rigorous structure. It is also considered by many developers as a relatively tedious phase. This project aims to study the requirements process itself, analyse its various constituent actions and properties, extract commonalities and propose aspects that may benefit from automated support, i.e. with as little human intervention and interpretation as possible. The use of automated tools for gathering requirements has become essential to the requirements-management process as it streamlines it, making it easier to track, test, analyse, visualise, and consequently to communicate to stakeholders [1].

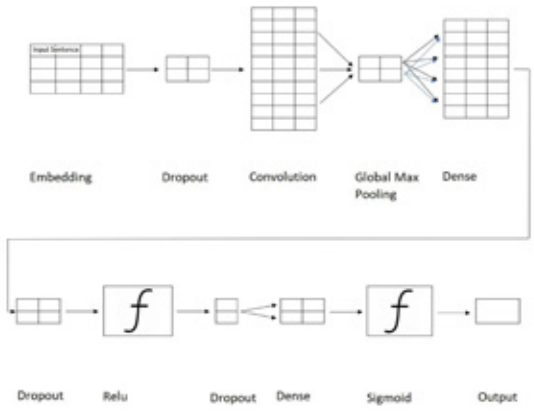

Figure 1. Structure of the CNN model

Requirements could be categorised into two types: high-level and low-level. Preliminary research was carried out in this area to obtain a good understanding of the categorisation of requirements by type. It was noted that when low-level requirements are well understood, this would facilitate the understanding of related high-level requirements [2]. However, the communicating of requirements and understanding between stakeholders could, at times, be subjective [3]. Therefore, research was also carried out regarding how such issues could be identified and managed, and if user experience might be a factor influencing this [4]. Further research was carried out concerning the manner in which human interaction could be mimicked in an automated fashion, and if any tools for this purpose exist. In order to gain a deeper understanding of any such tools, this study investigated the ways/methods in which requirements are elicited. Techniques involving computational aspects such as expert systems, fuzzy logic, machine learning applications and neural networks (alongside hash functions to reduce the data load) were subsequently considered when specifying the behaviour of the proposed solution. It was concluded that the most appropriate way to address the issue of automated requirements systems engineering would be through the use of a convolutional neural network (CNN). The requisite CNN, which acts as a binary classifier, could detect whether a requirement is 'good' or 'bad', according to a set criteria as to what constitutes a good requirement [5].

Like any other supervised machine learning design, the proposed solution could be divided into two parts: firstly, the training of the model, and secondly the testing (i.e. application) of the trained model. In simple terms, training is performed by passing the training data through the entire model, then comparing the predictions given by the model for the training data to the ground-truth labels. Once the model has finished training, the test data is then passed through. At this instance, the weights remain unchanged and the model would be evaluated on how well it managed to successfully classify the data. A dataset of raw system requirements was provided by KPMG. This dataset was then manually annotated according to a set of criteria that qualifies a requirement (i.e. indicates whether it is good or bad) [5].

The model performed very well on the test data, achieving high scores on metrics such as accuracy (0.90), precision (0.85), recall (0.88) and F-measure (0.86). When comparing the results to similar studies mentioned within the literature review, the model performed similarly and even better, in some cases. The manual checking of requirements is highly time-consuming and costly activity. The solution proposed within this study allows for the process to be carried out effectively and efficiently, with far less resources.

References/Bibliography:

[1] S. Aithal and P. Desai, An Approach towards Automation of Requirements Analysis.

[2] "Low-Level Requirements Software | Pi Innovo ProcessPi Innovo." https://www.pi-innovo.com/low- level-requirements-software/ (accessed Apr. 02, 2020).

[3] N. M. Azhari, "Conduct your Usability Testing As Soon As Possible (and How to Do it Properly) | Medium." https://medium.com/@nizarmaulana/conduct-your-usability-testing-as-soon-as-possible-and- how-to-do-it-properly-b61b643e6ee4.

[4] R. Thiayagarajan, "User Experience in Software Engineering," 2016, Accessed: Apr. 05, 2020. [Online]. Available: http://publications.lib.chalmers.se/records/fulltext/235397/235397.pdf.

[5] P. Zielczynski, Requirements management using IBM Rational RequisitePro. IBM Press, 2008.

GPU Implementation of Arithmetic for Very Large Integers

Gerard Tabone | **Supervisor: Prof. Johann A. Briffa**

Course: B.Sc. (Hons.) Computing Science

Arithmetic operations on very large integers are used in many applications such as computer algebra, computational number theory and public-key cryptography. A drawback is that most computers are typically built with a word size of 32 or 64 bits, which means that a machine register can only hold integers up to a maximum value of 2^64-1 for unsigned numbers. A solution would be to use multi-precision arithmetic libraries, which would allow numbers of any size. However, increasing the size of the numbers would also increase the amount of work required to perform the necessary arithmetic. Since a central processing unit (CPU) might struggle to keep up with the complexity of these computations, a graphics processing unit (GPU) could be used to offload some of the work from the CPU. GPUs and have been deemed to be useful in high-performance computing due to their substantial parallel architecture of hundreds to thousands of arithmetic units.

In this project, a parallel version of the addition, subtraction and multiplication operations for large numbers was implemented on a GPU. The addition was implemented using the prefix-parallel technique, which has been used on a smaller scale in the carry-lookahead adder [1]. This technique offers greater efficiency when executed on a large number of processors. For the multiplication, a warp-synchronous approach was adopted [2], taking advantage of how the thread grouping was done in the GPU. For multiplications of even larger numbers, the Karatsuba algorithm was chosen. The parallel implementation was tested on an NVIDIA GeForce GTX 1050 (Figure 1) and compared with the GNU Multiple Precision Arithmetic Library (GMP), which runs on a CPU. For numbers of 2^30 bits, a speed-up of 1.7 was achieved for the addition operation (Figure 2). Furthermore, the multiplication operation achieved up to 2.6 speed-up for 1024-bit numbers when fully utilising the GPU.

Figure 1. Similar NVIDIA GeForce GTX 1050 GPU used for testing (source: NVIDIA)

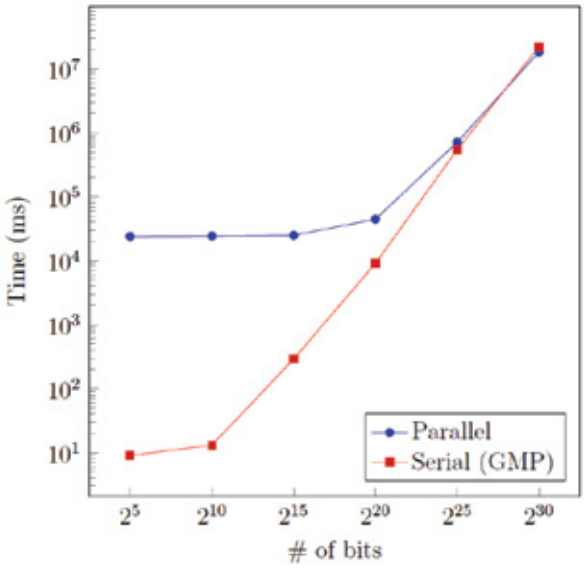

Figure 2. Timings for 1,000,000 additions comparing the serial and parallel implementations

References/Bibliography:

[1] A. Weinberger and J. Smith, "A one-microsecond adder using one-megacycle circuitry," IRE Transactions on Electronic Computers, no. 2, pp. 65-73, 1956.

[2] T. Honda, Y. Ito, and K. Nakano, "GPU-accelerated bulk execution of multiple-length multiplication with warp-synchronous programming technique," IEICE TRANSACTIONS on Information and Systems, vol. 99, no. 12, pp. 3004-3012, 2016.

Investigating real-time remote activity-scheduling to facilitate learning for children within the autism spectrum

Keith Vanhear **Supervisor: Dr Chris Porter**

Co-Supervisor: Prof. Paul Bartolo

Course: B.Sc. IT (Hons.) Computing and Business

The learning sciences have demonstrated a need for more flexible learning environments that encourage improved versatility in the educators' teaching approach, taking well into account the diverse needs of learners, thus enhancing their learning experience [1]. Flexible learning environments revolve around the learners' interests and diverse needs, making learning more personalised, accessible and engaging. One way of creating a flexible learning environment is through the use of digital technology [2].

Classrooms are becoming increasingly technology-rich environments [2] and new trends in personalised learning and access rely heavily on the use of digital technologies. These new developments have the potential of assisting educators when seeking to reach out to more diverse learners, including those with autism spectrum disorders (ASD) [3]. Literature regarding this topic has invariably highlighted that the prevalence of children diagnosed with ASD is soaring [4] [5] [6]. It is claimed that children with ASD encounter learning barriers due to their limitations in social interaction and communication [7] [8] [9]. In a technological age, assistive technology may help these students to improve their performance in their educational journey, and also in developing other life skills [10]. Literature in the field claims that digital technology has proven to be all the more effective and beneficial to education, since it increases engagement resulting from enhanced interaction with technology [11] [3].

This study focused on young children diagnosed with ASD, and their learning

environment. Various bodies of research have demonstrated that the use of activity schedules is effective in developing diverse skills in children with ASD, such as engagement in learning tasks, willingness for participation in activities, and independence in transitioning between tasks and/or situations [12] [13]. This study, which is exploratory by nature and adopts a mixed-method approach, seeks to demonstrate and reinforce the effective use of activity schedules and how they support students with ASD in their daily tasks.

Various interventions were made through an iterative approach, which enhanced the use of activity schedules as a digital tool. The design of the digital tool was informed by qualitative and quantitative techniques, including thematic analysis and standardised human-computer interaction (HCI) metrics involving various stakeholders, such as educators, learners and guardians. Furthermore, due to the unforeseen circumstances that emerged after a state of emergency was declared nationally in view of the COVID-19 pandemic, which restricted physical access to the classroom, this research evolved to adapt the use of activity schedules to a new reality of virtual learning. The emerging technology – code-named LivePlan – was developed as an online tool affording real-time collaboration facilities, as well as peer-to-peer voice and video communication, to enable educators, learners and their guardians to co-create and manage the activity schedules.

The main findings of this research suggest that even in an emergency such as the COVID-19 situation, because of which learning was taking place virtually, the proposed activity schedule real-time collaborative platform afforded leaners a smoother transition between tasks/activities, while contributing to create a virtual and flexible learning environment.

References/Bibliography:

[1] A. Meyer, D. Rose and D. Gordon, Universal Design for Learning: Theory and Practice, Wakefield, MA: CAST, 2014.

[2] OECD, Trends Shaping Education, Paris: OECD Publishing, 2016.

[3] K. Boser, M. Goodwin and S. Wayland, Technology tools for Students with Autism, Innovations that Enhance Independence and Learning, Baltimore, MD: Brookes Publishing, 2014.

[4] M. Haney, Understanding Children with Autism Spectrum Disorders: Educators Partnering with Families, Thousand Oaks, CA: Sage, 2013.

[5] Autism Europe, "Autism Europe," Autism-Europe aisbl, [Online]. Available: https://autismeurope.org. [Accessed 02 July 2020].

[6] Maenner, M.J.; Shaw, K.A.; Baio, J. et al, "Prevalence of Autism Spectrum Disorder among children aged 8 years - Autism & Developmental Disabilities Monitoring Network, 11 Sites, United States 2016," Centres for Disease Control and Prevention, Atlanta, GA, 2020.

[7] American Psychiatric Association, Diagnostic and Statistical Manual of Mental Disorders (5th Edition), Washington, DC: American Psychiatric Association, 2013.

[8] J. Arciuli and J. Brock, Communication in Autism., Amsterdam: John Benjamins Publishing, 2014.

[9] E. Kärnä, K. Dindar and X. Hu, "Educators' Engagement with Children with Autism Spectrum Disorder in a Learning Environment with Multiple Technologies in Finland and China," Interactive Learning Environments, pp. 1-15, 2018.

[10] S. Nur Syuhada Che Daud, M. Maria, F. Shahbodin and I. Ahmad, "Assistive Technology for Autism Spectrum Disorder: A Review of Literature," in Proceedings of International MEDLIT Conference, Kuala Lampur, 2018.

[11] V. Kirinić, V. Vidaček-Hainš and A. Kovačić, "Computers in Education of Children with Intellectual and Related Developmental Disorders," International Journal of Emerging Technologies in Learning, pp. 12-16, March 2010.

[12] R. Stromer, J. Kimball, E. Kinney and B. Taylor, "Activity Schedules, Computer Technology, and Teaching Children with Autism Spectrum Disorders.," Focus of Autism and Other Developmental Disabilities, vol. 21, pp. 14-24, 2006.

[13] L. E. McClannahan and P. J. Krantz, Activity Schedules for Children with Autism: Teaching Independent Behavior (2nd Edition), Bethesda, MD: Woodbine House, 2010.

The use of tablet application to develop the social interaction skills of children with intellectual disabilities

Luke Xuereb **Supervisor: Dr Peter Albert Xuereb**

Course: B.Sc. IT (Hons.) Computing and Business

With technology creeping into every aspect of everyday life and facilitating various daily tasks, education is no exception. Technology could, and should help, with improving the speed and quality of education. Teaching social skills is an area that tends to be somewhat overlooked in the education sector, meaning that when technology is implemented in education, it is generally oriented towards enhancing academic skills.

Social skills could be defined as the necessary behavioural components that are essential in correct interpersonal interaction, inevitably also averting the causing of physical or psychological harm to others [1]. Social-skills training (SST) has been derived from a number of approaches, all of which assume that social behaviour could be taught using a number of structured teaching methods. Some of these methods include homework tasks, discussion and modelling [1].

This project aims to apply the use of technology to the teaching of social skills to children with intellectual disabilities, focusing on the core skills required to behave within the social norms of society in a number of key social situations, such as the classroom, which forms a significant part of the daily lives of children.

This goal was achieved through the creation of an Android application, where, the children were provided with an innovative, interactive platform through which they could practice their social skills under the guidance of their learning-support educators (LSE). This is accomplished through a set of stories with a social theme, and with each 'social' story, the children are presented with three different illustrations. One of these illustrations portrays the correct way to behave, whereas the other two demonstrate incorrect behaviour. The children are encouraged to choose which illustration they think presents the correct behaviour. Gamification techniques, such as rewards and encouraging sound effects, were incorporated into the application to increase the children's level of engagement and motivation. The illustrations were displayed through the use of video-based instructions, which is the creation of videos by teachers to be used in a classroom environment, in such a way that the teacher would be in control of the exact, tailored content that is presented to each student [2]. Past researchers have proven that video-based instruction could be used successfully in helping adolescents with intellectual disabilities develop social skills [2].

Figure 1. A snippet of the story entitled, "Giving Good Advice"

The application was evaluated using a set of feedback forms, and surveys that were conducted among LSEs from two schools. More than one school was requested to participate in the evaluation of this study, so as to remove the possibility of any form of bias, which could have occurred if restricting the study to a single school.

The results obtained from the evaluation suggested that the proposed IT solution does increase a child's motivation and productivity in learning of social skills, and that gamification techniques play a crucial role in capturing children's attention, enabling them to perceive an obligatory task such as social skills training as an exciting activity which they look forward to. Furthermore, the results also indicated that such an IT solution could facilitate the teaching process of social skills from the perspective of an LSE. However, the results relating to this research question were not conclusive, and further research on this matter would be required.

References/Bibliography:

[1] S. Spence, "Teaching Social Skills to Children", Journal of Child Psychology and Psychiatry, vol. 24, no. 4, pp. 621-627, 1983, Available: 10.1111/j.1469-7610.1983.tb00138.x.
and the second to the following:

[2] J. Plavnick, T. Kaid and M. MacFarland, "Effects of a School-Based Social Skills Training Program for Adolescents with Autism Spectrum Disorder and Intellectual Disability", Journal of Autism and Developmental Disorders, vol. 45, no. 9, pp. 2674-2690, 2015. Available: 10.1007/s10803-015-2434-5.

Implementation of Synchronised Multi Camera Video Capture on an Embedded System

Kyle Borg **Supervisor: Prof. Johann A. Briffa**

Course: B.Sc. (Hons.) Computer Engineering

A light field array consists of multiple cameras that capture different views of the same scene at the same time. The frames captured can be merged together to form a basic 3D model of the captured object. Frame synchronization is necessary to ensure that all the cameras capture the scene at the same instant.

The aim of this project was to implement a multi camera system based on an NVIDIA Jetson TX2 microcontroller. The CMOS cameras used were from Leopard Imaging, where the MIPI interface was used to trigger the cameras. The available cameras and adapter board do not support frame synchronization in hardware, so the aim of this thesis was to minimise the timing discrepancy between the cameras as much as possible.

Several tests were conducted to compare alternative ways to trigger the captures, to ensure that the final system is functioning correctly, and to determine the synchronization error between the frames obtained from each camera. Furthermore, a user interface with physical buttons was implemented, making the system portable and user friendly. This was implemented by making use of the GPIO module provided on the development kit of the microcontroller.

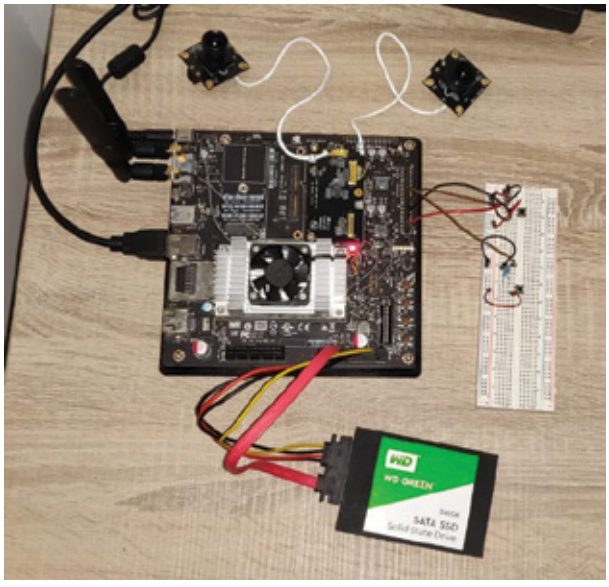

Figure 1. The hardware used for this project

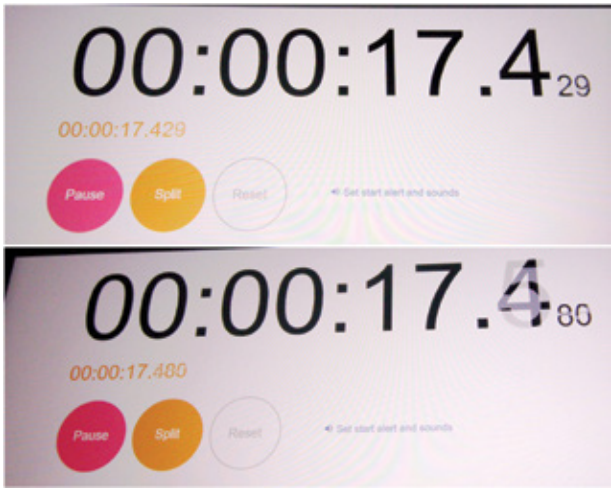

Figure 2. Two frames captured that are comparing the discrepancy between the cameras using a digital stopwatch

The tests were performed by capturing frames of two types of clock animations. First, images of an analogue clock were captured, where the discrepancy between the frames was indicated by the difference in the position of the second hand. The smaller the difference between the position of the second hand, the smaller the synchronization error, hence a better synchronization between the camera modules. Another set of tests were conducted by capturing frames of a digital stopwatch. In this case, the synchronization error between the cameras could be obtained by calculating the difference between the times recorded in the frames. Again, the smaller the difference in time between the captured frames, the better the synchronization. Furthermore, when possible, the software also tracked the kernel timestamps of the captured frames. The difference between the recorded timestamps is equal to the synchronization error between the cameras. Each test was performed multiple times (up to ten times) and an average synchronization error calculated.

References/Bibliography:

[1] "Light fields and computational photography", Graphics.stanford.edu, 2020. [Online]. Available: http://graphics.stanford.edu/projects/lightfield/.

Development of an FPGA Based ECU for a Petrol Engine

Andrew Muscat **Supervisor: Prof. Ivan Grech**

Co-Supervisor: Prof. Mario Farrugia

Course: B.Sc. (Hons.) Computer Engineering

All modern automotive engines feature an engine control unit (ECU), which is a device for controlling emissions, performance, and efficiency. An ECU is activated through a number of sensors within an engine that constantly monitor it within set parameters and, from the information obtained, the ECU conditions the engine to control certain variables, such as amount of fuel sent into the engine and when to start it.

Many automotive companies strive to employ an ECU that could ensure that the engine would run to its full potential. Some ECUs feature different engine mappings, making it possible to switch from 'economy mode' to 'sport mode', thus switching from a fuel-efficient mode to a more powerful performance. However, certain functions could only be achieved in conjunction with an ECU. One example is the Volkswagen active-cylinder technology (ACT), which is designed to automatically stop the fuel supply to two of the engine cylinders when carrying a lighter load, in order to save fuel. Such functions make it possible to continue producing higher-performance engines, paying due attention to environmental obligations, by providing an optimal balance between performance, fuel economy and lower emissions.

This project consists of developing the hardware for an ECU based on a Genesys Virtex 5 FPGA platform, together with the necessary electronic interfaces required for monitoring sensors and driving actuators. The programming was carried out using the hardware-description language VHDL in order to define the digital module used to process the information received from sensors, and outputting the signals to the engine components. The electronic circuitry was employed to interface the crank position sensor, MAP sensor, and spark plugs, and for driving the injectors. The TLP250 optocoupler was especially used to isolate and protect the FPGA (field-programmable gate array) from potentially dangerous voltages, as well as to drive higher-voltage components.

In order to implement and test the system safely, each module was built and tested separately. Components such as the injector, spark plug, and CPS modules were built and tested with hard-coded HDL blocks. Once their operation was confirmed, they were connected to their pertaining PMOD ports and circuit blocks, and tested again with an oscilloscope to confirm the voltage output. The timing of these components was also tested with a number of output flags to confirm that the various signals were switching on at the right time, and for the right length of time.

The final product contained three main hardware components, namely: an FPGA, a breadboard circuit for the interface electronics, and the several actuators and sensors on the vehicle itself. Each of these vehicle components was tested from the circuit output and FPGA, and it was confirmed that the FPGA could properly drive the multiple actuators and detect all of the sensors.

Since accuracy was a major consideration in the project, it was necessary to test the timing constraints thoroughly. To this end, the clock speed of the FPGA was set to 10 nanoseconds and the timing accuracy was to 10 nanoseconds. This was deemed highly satisfactory, and in excess of what would be required from an automotive engine. All the basic functions were tested and confirmed to work properly with the given timing constraints down to the 10- nanosecond clock.

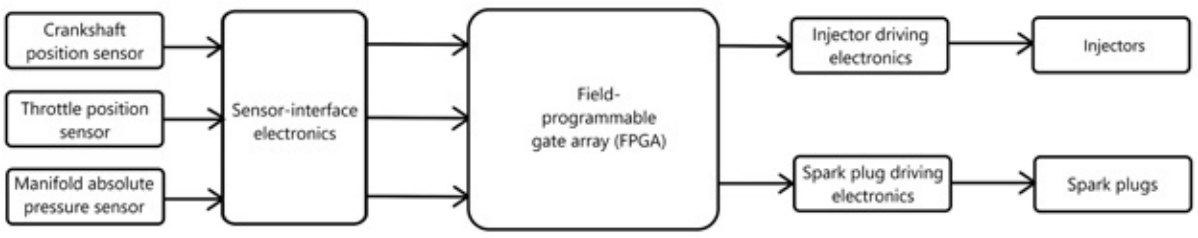

Figure 1. Block diagram of the system and actual implementation

An IoT-Based Forest Fire Detection System

Dillon Scicluna **Supervisor: Dr Gianluca Valentino**

Course: B.Sc. (Hons.) Computer Engineering

The objective of this project was to design an IoT-based (internet of things) system that could detect fire as quickly as possible, in order to alert authorities before it spreads over an extensive area. The proposed system consists of an Arduino board integrated with a temperature, humidity and pressure sensor that communicates via long-range IoT protocols, such as LoRaWAN, to relay information about their environment to a central server. Readings would be measured and sent wirelessly every couple of seconds. These are saved in a specified text file, from which graphs measuring each of the three weather variables (temperature, pressure and humidity) against time respectively could be generated to detect the change of rate of that particular variable being monitored.

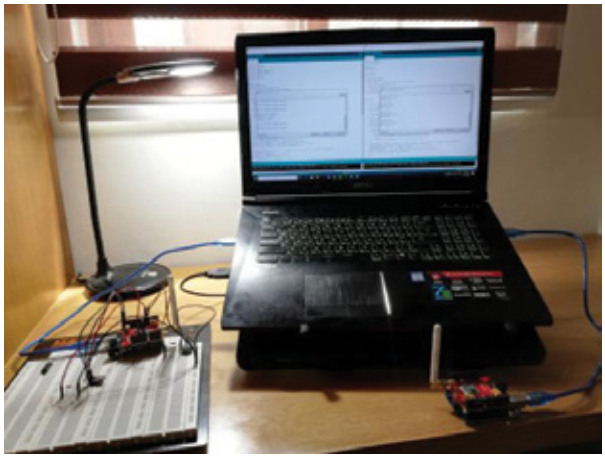

Figure 1. Set-up of the system

The process outlined above would be necessary, as setting a threshold would not be enough to detect whether there is fire or simply a change in climate. The intention of using a multi-sensorial approach is to reduce false alarms and provide more accurate monitoring of the area under surveillance. Such a system would reduce the time needed to detect outdoor fire, resulting in less damage to the area. A key outcome observed from the results obtained was that the sensor module could detect changes in all three weather variables even up to a distance of 7.5 metres from the heat source – and potentially even farther away – obeying the laws of physics and the inverse-square law, which states that temperature and pressure increase whilst humidity decreases. The simulations were performed with a successful wireless transmission from one microcontroller to the other, both connected to different power supplies and located in different areas from each other to test the LoRaWAN protocol-connectivity distance.

In conclusion, it could be deduced that from all the simulations carried out, the system performed to expectations and produced the anticipated positive outcomes. In other words, the system gathered measurements from the sensor module and saved these as a text file for records purposes and the generation of graphs, and sent the data via IoT to a microcontroller acting as a base station situated at a distance from the area under observation.

Figure 2. A trial of the system, deployed in the Buskett woodland, in Malta

Developing an educational game for children with speech sound disorder

Kurt Camilleri **Supervisor: Dr Lalit Garg**

Co-Supervisor: Dr Peter Albert Xuereb

Course: B.Sc. IT (Hons.) Software Development

Speech sound disorder (SSD) is a broad term covering all the disorders, as classified by Dodd [1]. The omission or substitution of a phoneme, which is a distinct sound that distinguishes one word from another, is a common occurrence in persons with SSD. Children especially may find a certain speech sound difficult to express not only in a particular word but in any occurrence of that precise sound.

Much progress has been made on the subject of SSD treatment techniques such as the cycles approach and the minimal-pairs approach [2,3]. These advancements could be supported by technology in order to implement the techniques as affectively as possible, while also increasing the child's motivation to practise. The involvement of technology in children's learning has proved to be useful and effective when used correctly. Indeed, it has been noted that pupils making use of technological aids were very co-operative and active in the classroom [4]-[6]. Children and teenagers today are among the largest group of tablet and smartphone users, and it is had been predicted that children between the ages of 3and 4 will use mobile devices more extensively during their lives, than their parents would [7]. Furthermore, these studies show how mobile games are one of the main reasons why children use tablets and smartphones on a daily basis because of their evident appeal [7].

This study followed the development of an educational game to bridge the gap between technology, and speech and sound learning. The game was developed on both iOS and Android platforms and was tested on multiple devices, including tablets. The pupils' native language was adopted in testing the functionality of the devices used for transforming the speech produced by the user into text. The written word resulting from the user's speech was then compared with the word displayed. The game implements treatment techniques such as minimal pairs and the cycles approach.

Figure 1. Screenshots of the educational game

Furthermore, it incorporates game development methods, such as providing constant feedback to the user and using a bright colour scheme [8]. The game also produced statistics for speech and language therapists in order to identify five phonemes with which the child was having the most difficulty.

This study evaluates the effectiveness of an educational game in the field of speech production through the iterative process, and the evaluation meetings with speech and language therapists. This study also sought to assess the ability of the native-speech-to-text technology found in most mobile devices to be used for children with SSD. Qualitative data was gathered with the aid of a questionnaire, which provided two key findings. The first finding suggests that the use of an educational game was successful in motivating the children concerned in participating and to practise their speech. The second suggests that the native-speech-to-text technology could be quite difficult for a child with SSD to use, and therefore this technology could be further improved upon to be more effective for children with SSD.

References/Bibliography:

[1] Dodd, B. (2014). Differential Diagnosis of Pediatric Speech Sound Disorder. Current Developmental Disorders Reports, 1(3), 189–196. https://doi.org/10.1007/s40474-014-0017-3

[2] Prezas RF, H. B. (210 C.E.). The cycles phonological remediation approach: Enhancing children's phonological systems (M. R. Williams L, McLeod S (ed.)). Baltimore, MD: Brookes Publishing Company. pp. 137-158.

[3] Weiner, F. (1981). Treatment of phonological disability using the method of meaningful minimal contrast: two case studies. The Journal of Speech and Hearing Disorders, 46(1), 97– 103. https://doi.org/10.1044/jshd.4601.97

[4] Diemer, T. T., Fernandez, E., & Streepey, J. W. (2012). Student Perceptions of Classroom Engagement and Learning using iPads. Journal of Teaching and Learning with Technology, 1(2), 13–25.

[5] Chen, P. S. D., Lambert, A. D., & Guidry, K. R. (2010). Engaging online learners: The impact of Web-based learning technology on college student engagement. Computers and Education, 54(4), 1222–1232. https://doi.org/10.1016/j.compedu.2009.11.008

[6] Wario, R. D., Ireri, B. N., & De Wet, L. (2016). An Evaluation of iPad as a Learning Tool in Higher Education within a Rural Catchment: A Case Study at a South African University. Proceeding International Conferences ITS, ICEduTech and STE 2016 AN, 9. http://login.ezproxy1.lib.asu.edu/login?url=https://search.proquest.com/docview/1968427998? accountid=4485

[7] Markov, M., & Grigoriev, Y. G. (2013). Wi-Fi technology-an uncontrolled global experiment on the health of mankind. Electromagnetic Biology and Medicine, 32(2), 200–208. https://doi.org/10.3109/15368378.2013.776430

[8] Gee, J. P. (2008). Learning and games. The Ecology of Games: Connecting Youth, Games, and Learning, 21–40. https://doi.org/10.1162/dmal.9780262693646.021

Learning Models using Similarity Based and One vs Previous Paradigms

Daniel Cauchi | **Supervisor: Prof. Ing. Adrian Muscat**

Course: B.Sc. (Hons.) Computing Science

Classification between two classes of objects is a common-use case for the creation of machine learning models. The general practice is to create a model consisting of a single binary classifier that could discriminate between the classes by putting the examples of the first class on one side of a generated decision boundary, and the examples of the second class on the other. However, it is often the case where it would be necessary to discriminate between more than two classes, which is a problem known as 'multi-class classification' [1].

A popular strategy adopted by multiple models is to use an ensemble of binary classifiers, built using some paradigm. These would collectively be able to discriminate between the series of classes. The paradigm, which would determine how the binary classifiers are built and interact with one another, is the focus of this project. Traditional paradigms have some notable limitations, such as scaling up poorly to very large datasets, necessitating the retraining of the entire model from scratch when a small update would be required. These updates may consist of adding more data related to current classes, or even including new classes, thus requiring an increase in classifiers. A prominent paradigm that harbours these limitations is One vs Rest [2], which builds a classifier for each class, such that each classifier uses all the examples from every class. Since the time taken to train each classifier would depend on the total number of training examples, this paradigm would take a long time to train on large datasets.

In an effort to overcome some of these limitations, this study presents three novel paradigms for multi-class classification. The starting point is Similarity Based Learning (SBL), followed by the One vs Previous (OvP) method (Diagram (a) in the accompanying image). Lastly, these two would be combined into a modified

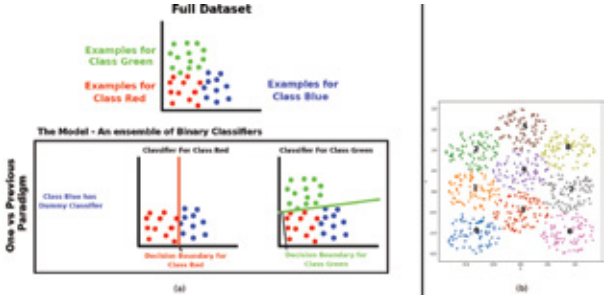

Figure 1. (a) Overview of One vs Previous Paradigm (b) Visualisation of synthetic dataset

algorithm of SBL (OvPSBL). The novel paradigms attempt to minimise the issue of poor scaling by decreasing the number of training examples required by each classifier. Moreover, in the event of new data entering the system, retraining from scratch could be avoided by decoupling the classifiers. This would require the retraining of certain classifiers only, rather than the entire model, thus making the novel paradigms suitable in both normal classification scenarios as well as lifelong machine learning [3, 4].

In this study, three paradigms have been evaluated over a purposely defined synthetic dataset (Diagram (b) in the accompanying image) as well as a real-world dataset (SpatialVOC2k) [5]., The presented paradigms were compared to a baseline paradigm so as to demonstrate how these scale up when changing the dataset characteristics, in relation to each other and traditional methods. Results show that, under most conditions, the studied paradigms are more efficient than current methods, whilst retaining a good classification accuracy. Furthermore, different paradigms work better than others in different situations. While this project offers a preliminary outline of how the paradigms perform, future work would help uncover further improvements and shed light on scenarios where certain paradigms would be more effective than others.

References/Bibliography:

[1] M. Aly, "Survey on multiclass classification methods," in Technical Report, Caltech, 2005.

[2] R. Rifkin and A. Klautau, "In defense of one-vs-all classification," Journal of Machine Learning Research, vol. 5, pp. 101-141, 12 2004.

[3] D. Silver, Q. Yang, and L. Li, "Lifelong machine learning systems: Beyond learning algorithms," in AAAI Spring Symposium: Lifelong Machine Learning, 03 2013.

[4] A. Blum, "On-line algorithms in machine learning," in Developments from a June 1996 Seminar on Online Algorithms: The State of the Art, (Berlin, Heidelberg), pp. 306-325, Springer-Verlag, 1998.

[5] A. Belz, A. Muscat, P. Anguill, M. Sow, G. Vincent, and Y. Zinessabah, "Spatialvoc2k: A multilingual dataset of images with annotations and features for spatial relations between objects," in Proceedings of the 11th International Conference on Natural Language Generation, (Tilburg University, The Netherlands), pp. 140"145, Association for Computational Linguistics, Nov. 2018

Implementation of an electric guitar pedal on an FPGA

Peter Galea St John **Supervisor: Dr Ing. Owen Casha**

Course: B.Sc. (Hons.) Computer Engineering

Implementing real-time digital signal processing could be computationally expensive when developed as software to be run on a generic computer platform. A solution to compensate for the high computational demand would be to implement it as an application-specific hardware. This could be achieved in various ways, one of which being through field-programmable gate array (FPGA), which permits an inherent parallelised architecture that facilitates real-time processing.

The aim of this project was to develop and implement different digital audio effects, using a hardware definition language. In particular, this work presents the implementation of a guitar-effects pedal board, based on an FPGA. A block diagram representation of the implemented system is shown in Figure 1. The Atlys Spartan-6 FPGA trainer board was used since it features an AC-97 audio codec with line-in/line-out, microphone and headphone interfaces, which facilitated the implementation. A number of effects were introduced, including: echo, reverb, chorus, distortion, tremolo, phaser and flanging. The implementation of these effects was aided by considering the Z-transform transfer function of each effect. The implementation was designed around the AC-97 audio-codec hardware driver developed by Tony Storey, and obtained from Digi-Key. This driver runs at a 48 kHz sampling rate and features an 18-bit signal representation. An analogue signal-conditioning amplifier, based on an OPA 350 operational amplifier, was also designed and implemented on a printed circuit board to adequately amplify the signal from an electrical guitar

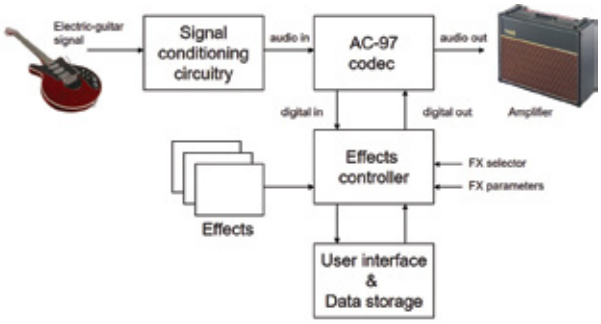

Figure 1. Block diagram of the implemented guitar-effects pedal board

prior to feeding it into to the ADC of FPGA trainer board. It is to be noted that the AC-97 audio codec has its own frequency-dependent phase-shift contribution which had to be taken into consideration while carrying out the various tests.

The selected effect was applied to the raw digitised signal, which was eventually converted to the analogue domain through the on-board DAC, and outputted to an amplified speaker. The performance of the implemented system was tested and verified while comparing the measured results with those obtained from the ideal implementations carried out in MATLAB. Depending on the nature of the effect, the frequency response and/or the transient response of each effect was verified by means of a high-speed digital oscilloscope and a signal generator, as shown in Figure 2.

Sign Language Sensorial Glove

Lauri Anastasia **Supervisor: Prof. Ivan Grech**

Co-Supervisor: Prof. Ing. Victor Buttigieg

Course: B.Sc. (Hons.) Computer Engineering

This project aims at developing a sensor glove that could process hand gestures and translate them into the corresponding alphabetical letter according to the American Sign Language (ASL). This would allow a deaf person to communicate using the said sign language alphabet to convey a message, without requiring the person at the receiving end to be conversant with ASL.

The glove incorporates a number of sensors sensitive to the smallest degree of the bending of each

finger. Flex sensors were the optimum choice for this function in view of their long lifespan, durability and ease of integration, given that their resistance changes in proportion to the bending of the sensor [1]. Ten flex sensors were used for this project to ensure that any bending in the two main joints of the fingers could be detected.

Figure 1. Image of Sensor Glove

The glove would also be required to detect hand orientation. For this purpose, the BNO055 sensor was used. This enables the system to distinguish between two letters that employ the same finger formation, but with a different orientation. The same sensor was also used to detect any linear acceleration caused by the signing of an arch, which represents the repetition of a letter. The BNO055 was chosen since it interfaces seamlessly with the LilyPad Arduino Main Board microcontroller. Using the appropriate

Figure 2. System Block Diagram

software libraries, the BNO055 would provide the necessary orientation and linear acceleration values, which are then processed by the LilyPad [2].

The code was executed on the LilyPad microcontroller and translated the sensory data into the corresponding alphabet character, which was output on the Arduino IDE console. The data transmission from the LilyPad to the PC was handled by the HC-06 Bluetooth module, whereas power was provided by an armband-mounted battery, thus ensuring wireless operation.

In order to test the system accuracy, the glove was used to display each of the twenty-six letters for ten times. Nineteen out of the twenty-six letters were displayed correctly without any delay or false reading.

Three letters – 'C', 'P' and 'W' – were incorrectly detected 10% of the time, whereas the letter 'J' was incorrectly output as the letter 'I' four out of ten times. There were also nine instances where the output letter ('I', 'U', 'V' and 'W') was delayed by a few seconds until the fingers attained the predetermined formation. Most of these issues could be attributed to slight movements of the flex sensors relative to the glove fabric, resulting in variations in the angle of the bend even when the same letter was being signed.

References/Bibliography:

[1] "Flex Sensor Datasheet," SpectraSymbol, [Online]. Available: https://www.sparkfun.com/datasheets/Sensors/Flex/flex22.pdf. [Accessed 12 December 2019].

[2] Bosch, "BNO055 Intelligent 9-axis absolute orientation sensor," Bosch Sensortec, November 2014. [Online]. Available: https://cdn-shop.adafruit.com/datasheets/BST_BNO055_DS000_12.pdf . [Accessed November 2019].

Classification of Deceptive Traits from Audio-Visual Data

Braden Refalo **Supervisor: Dr Andrea De Marco**

Co-Supervisor: Dr Claudia Borg

Course: B.Sc. IT (Hons.) Artificial Intelligence

The automated analysis and inference of human behavioural traits by machine interfaces is a growing field of research. On the other hand, research regarding the study of deceptive traits and classification is relatively scarce. It has been established by various studies [1, 3] that the augmentation of multi-modal information (e.g. acoustic analysis of speech and visual lip-reading) could enhance the performance of speech recognition systems traditionally geared exclusively towards acoustic data. Similarly, the analysis of speech augmented with human body language and facial patterns could help provide information on traits such as emotional state, or whether a speaker is trying to deceive an audience or interlocutor.

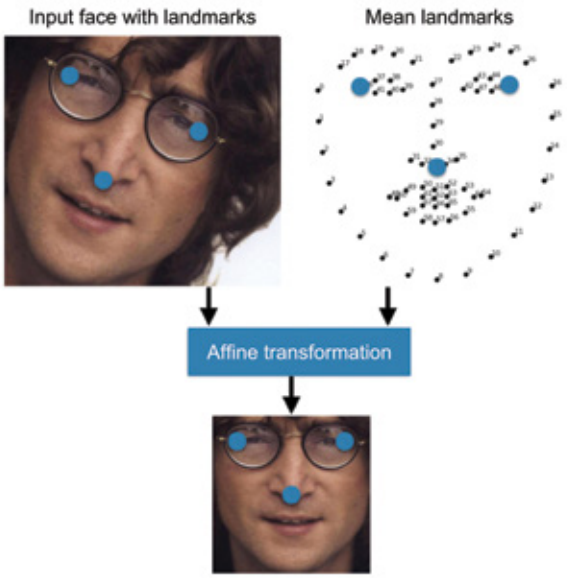

Figure 1. OpenFace visualisation for feature extraction [2]

Deception is an intriguing concept due to the fact that it is often misconceived by humans. As with sarcasm, understanding deception requires experience, yet there is no absolute and objective metric to determine whether someone is being deceptive. Sen et al. [1] developed an automated dyadic data recorder (ADDR) framework,

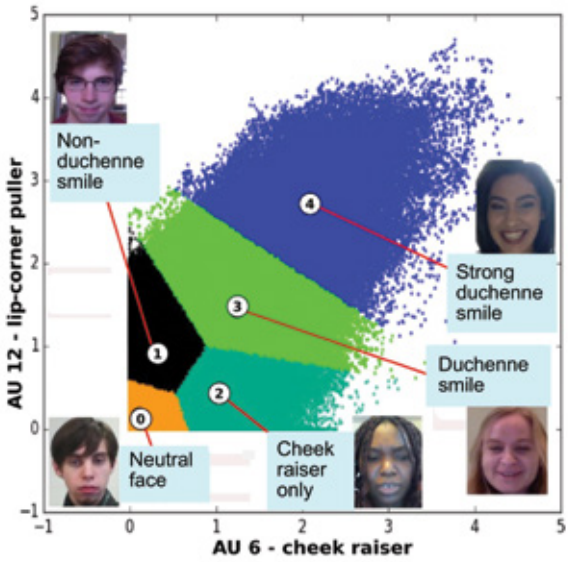

Figure 2. Segmenting facial expressions from OpenFace action units [2]

which led to the first comprehensive deception database composed of audio-visual interviews between witnesses and interrogators. Although this was not the first work in the field, previous datasets were either not accessible or not optimised for machine learning (ML) techniques and tasks [1].

This project investigates various applications of ML techniques in attempting to detect deceptive actions and encapsulate those traits. Various ML models and hyperparameters were explored within the hyperparameter space with Bayesian optimisation tuning. The utilisation of the best performers played a very important role in producing the final classifier. The proposed system is composed of a deep neural network with recurrent long-short-term memory (LSTM) cells. It is fed with aligned audio and OpenFace [2] features.

This study shows that a machine can capture and analyse deceptive traits at an adequate degree of confidence whilst achieving an AUC ROC score of 0.607. The existence of such an ML classifier suggests that there exists some pattern or latent model that defines deceptive actions.

References/Bibliography:

[1] T. Sen, M. K. Hasan, Z. Teicher, and M. E. Hoque, "Automated dyadic data recorder (ADDR) framework and analysis of facial cues in deceptive communication," 2017.

[2] Brandon Amos, Bartosz Ludwiczuk, and Mahadev Satyanarayanan. Openface: A general-purpose face recognition library with mobile applications. Technical report, CMU-CS-16-118, CMU School of Computer Science, 2016.

[3] T. Sen, M. Hasan, M. Tran, Y. Yang, and E. Hoque, "Say cheese: Common human emotional expression set encoder and its application to analyze deceptive communication," pp. 357–364, 05 2018.

Automated Report Generation from Football Match Commentary

Jake Seracino **Supervisor: Dr Joel Azzopardi**

Co-Supervisor: Mr Nicholas Mamo

Course: B.Sc. IT (Hons.) Artificial Intelligence

The sheer popularity of football means that most matches are extensively covered before, as well as after, the match. One particular practice that most news outlets follow is that of writing a brief, post-game report on the main highlights occurring throughout the match, such as goals scored or controversial decisions taken by the referee. Needless to say, such a task tends to be time-consuming, as it requires the writer to watch the match and then write the report. Moreover, online portals seek to produce and upload the report in as little time as possible from the end of the match, while the interest is still at its peak.

62'	Sergio Aguero (Manchester City) produces a lovely chip pass into the box, but one of the opposition players cuts it out. The ball goes behind for a corner. Manchester City will have an opportunity to threaten the opposition's goal.
61'	Raheem Sterling (Manchester City) receives a diagonal pass, bursts into the box and shoots towards the bottom left corner. Mathew Ryan is well positioned, though, and it's a comfortable save.
59'	Oleksandr Zinchenko (Manchester City) commits a rough challenge and Michael Oliver blows his whistle for a foul.
57'	Ilkay Gundogan (Manchester City) takes a first-time shot from the edge of the box, but his own player gets in the way and blocks it.
56'	Mathew Ryan displays his ability to read the game and clears the danger after Raheem Sterling (Manchester City) threatened to latch onto a through ball from just outside the box.
54'	Lewis Dunk (Brighton) reacts well and displays good awareness near the penalty spot to get on the end of a cross from the free kick. He jumps high and produces a header which goes well over the crossbar.
54'	Anthony Knockaert (Brighton) will take the free kick.
53'	Kyle Walker (Manchester City) manages to dispossess his opponent of the ball, but he commits a foul in the process and the referee blows his whistle. Brighton are awarded a free kick. Let's see what they create from this.
51'	David Silva (Manchester City) commits a foul and Michael Oliver immediately signals a free kick.

Figure 1. A football commentary transcript (extract)

The main aim of this research was to propose a suitable architecture to automatically generate a football-match report from the commentary of any given match. The problem was framed as a special kind of extractive summarisation, whereby each sentence in the commentary was considered as a candidate for inclusion in the final report. From each 'candidate sentence', a number of features were extracted so that they could be scored and ordered.

Manchester City came from behind to retain the Premier League title in style with a 4-1 victory at Brighton on the final day of the season.

City looked on course for a nerve-jangling final outing at the Amex Stadium when Sadio Mane's opener for Liverpool against Wolves - which temporarily moved the Reds top of the table - was compounded by Glenn Murray heading Brighton into the lead on 27 minutes.

Figure 2. An extract from a manually written match report

The features employed for this study include those found in typical summarisation tasks, such as the position of a candidate sentence, as well as its sum of TF.IDF weights. A number of domain specific features, such as the inclusion of explicit highlight markers, were also utilised. These features are mainly based upon others used in similar research. After the commentary, the sentences were organised in the best combination of sentences, according to a set word count, and were subsequently selected to compile the final report.

The study has taken into account that certain events occurring during a given football match tend to occur frequently together and in an ordered manner. This research has defined such phenomena as episodes and recognises them using sequential rule mining. The study also proposes a secondary architecture, which also accounts for episodes when compiling the final report. The performance of this data pipeline is compared and contrasted with one that does not consider episodes. Finally, the performance of both pipelines was measured against a number of task-specific and general summarisation baselines to evaluate their validity in addressing the issue under discussion.

A Machine Learning Approach to Predict Epileptic Seizures from EEG Data

Jonathan Cauchi | **Supervisor: Dr Lalit Garg**

Course: B.Sc. (Hons.) Computing Science

In light of the recent progress in machine learning (ML) and artificial intelligence (AI), researchers aim to apply techniques for improving and automating certain facets of clinical practice. One of the more intriguing and compelling applications of modern computing in a healthcare context is the early detection and prediction of disease. In the case of epilepsy, the prediction of seizure onsets would allow patients to adequately prepare for such recurrent episodes of convulsion, thus improving their quality of life. Although seizures are preventable by specific medication and therapies, it is not uncommon for patients to suffer from intractable seizures caused by drug-resistant epilepsy. The prediction of seizure onsets would allow patients some relief in knowing when to be prepared and when to avoid activities that could be potentially dangerous, such as driving.

Classifier	Accuracy	Specificity	Sensitivity
kNN	66.93	69.39	64.47
SVM	69.07	69.80	68.34
Ensemble	74.59	77.30	71.89

Figure 1. The receiving operative characteristic (ROC) curve

This dissertation presents a review of the performance of a set of supervised machine learning methods for the task of seizure prediction. The study involves using a dataset that includes non-invasive scalp electroencephalography (EEG) signals, which are brain electrophysiological readings that do not involve surgery. The findings are presented

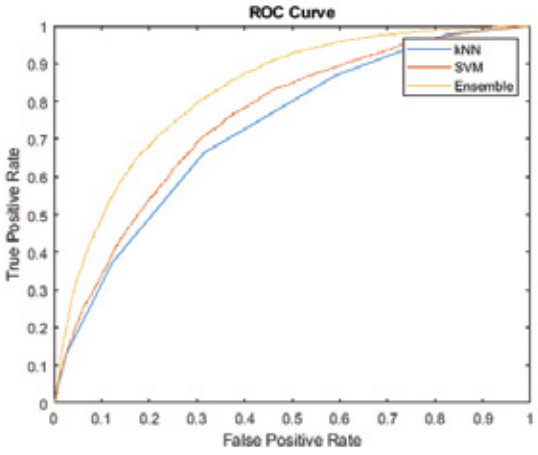

Figure 2. Performance metrics

and compared with current research in the field. Upon completing data labelling and pre-processing, statistical and wavelet features from the signals were subsequently extracted, which in turn was used for training the machine learning models. The results obtained from the k-nearest neighbour (kNN), support vector machines (SVM) and ensemble classifier were compared accordingly. The results have been reported on the CHB-MIT dataset, which includes 198 seizure readings from 22 patients suffering from intractable seizures. The study suggests that the three above-mentioned methods deliver a similar level of efficiency. However, the ensemble classifier achieved a higher specificity, sensitivity and accuracy.

References/Bibliography:

[1] Shoeb, Ali & Guttag, John. (2010). Application of Machine Learning To Epileptic Seizure Detection. ICML 2010 - Proceedings, 27th International Conference on Machine Learning. 975- 982.

[2] (4 September, 2017). Epilepsy - Treatment - NHS. Available: h ttps://www.nhs.uk/conditions/ e pilepsy/treatment/.

NutriPeer: Design and implementation to help reduce nutrient misconsumption

Abigail Cortis | **Supervisor: Dr Lalit Garg**

Course: B.Sc. IT (Hons.) Software Development

Most of the persons who download a health-related mobile application tend to use such applications for a few months at most, before losing interest in it [1]. This project aimed to help find design principles and features that would help increase the use of health-related mobile applications. To achieve this, an application – NutriPeer – was created specifically for this project.

One of the main features of NutriPeer is a recipe-sharing platform, where users could share their recipes with other users, and therefore the pattern of offering the same advice, is broken. Another method to try to increase the duration of the interest in using the mobile application was by performing a survey before designing the application to determine what prospective users would like a nutrition-oriented mobile application to include, but also to facilitate usability testing once the application was developed.

During the development phase, it was ensured that the design principles in Material Design [2],along with coding best-practices, were duly observed. At the requirements gathering stage, a survey was conducted asking potential users about which features they would like to see in a nutrition-oriented application that would motivate them to use the application, since no participant was using a nutrition-oriented mobile application. The most sought-after feature was a barcode scanner to log food consumption, given the perceived ease of use of such a feature. Another request was to include a recipe-sharing function, through which users could post and share their recipes, allowing other users to rate them according to their view of the recipe. Subsequently, usability testing was performed to ensure that the application would meet user expectations when seeking to use a nutrition-oriented mobile application. This was conducted by giving participants access to NutriPeer for a pre-set period of time, at the end of which user feedback was sought through a set of questions.

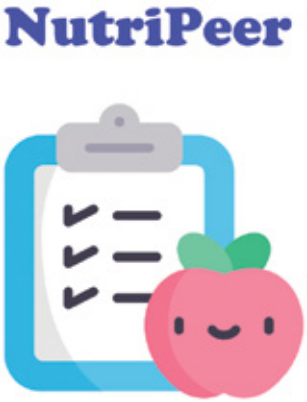

Figure 1. The NutriPeer logo

The application developed for this project includes the function to keep a log of food items that the user has consumed during the day, while giving an indication of the required daily intake of nutrients such as carbohydrates, protein, fibre and fat, as well as offering recipes to help the user achieve the recommended nutrient intake as suggested by the World Health Organization (WHO) [3]. Due to a high demand in the requirements gathering phase, the application also includes a water-intake diary to facilitate keeping track of the amount of water users would have consumed during the day.

References/Bibliography:

[1] K. Anderson, O. Burford, and L. Emmerton, 'Mobile health apps to facilitate self-care: A qualitative study of user experiences', PLoS One, vol. 11, no. 5, pp. 1–21, 2016.

[2] 'Design - Material Design'. [Online]. Available: https://material.io/design. [Accessed: 25-Aug-2020].

[3] 'WHO | 5. Population nutrient intake goals for preventing diet-related chronic diseases', WHO, 2007. [Online]. Available: https://www.who.int/nutrition/topics/5_population_nutrient/en/#:~:text=1 Background,value" for such a goal. [Accessed: 25-Feb-2020].

A study on the effectiveness, robustness and safety of smart wheelchairs in well-defined spaces

Dejan Mijailovic **Supervisor: Dr Michel Camilleri**

Co-Supervisor: Dr Peter Albert Xuereb

Course: B.Sc. IT (Hons.) Software Development

Transportation of people and objects in indoor areas requires several costly resources. Different environments such as a warehouse, workplace, or hospital, have similar requirements and face similar problems. In specific, individuals using a wheelchair who require another person's help to move from point A to point B seek facilitation to make their lives easier. Smart wheelchairs could have the potential to provide a solution to some of their needs. A smart wheelchair is an upgrade to a powered wheelchair and a normal wheelchair as it merges the use of artificial intelligence and mobile robotics with the present Power Wheelchair technology.

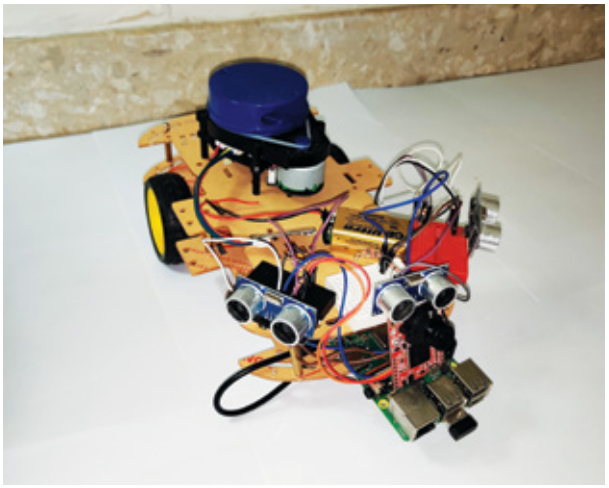

Figure 1. Prototype of the robotic vehicle

By adding an extra level of intelligence, smart wheelchairs may overcome many of a normal wheelchair's limitations by performing autonomous movement with minimal user intervention. Moreover, on another level, the wheelchair may be made capable of navigating a building autonomously and safely by exploiting various sensor technologies and computer-based logic. A variety of prototypes have been developed by several global research institutions. Although these mobility systems are still being developed and tested, they have only recently been introduced to the commercial world due to their high costs.

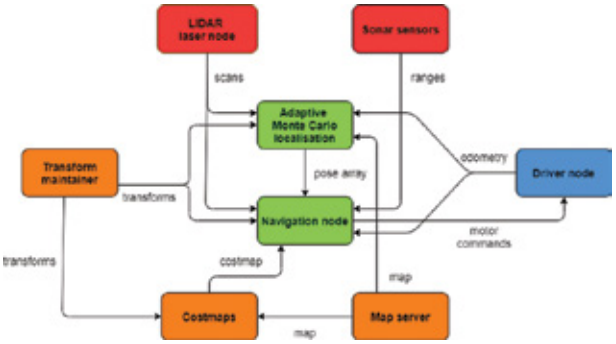

Figure 2. ROS node functionality diagram

The aim of this study is to replicate a smart wheelchair through a robotic prototype with the use of the low cost and readily available hardware, and to measure its effectiveness robustness and safety for the intended task in an objective and methodical manner. The prototype was installed with a motor driver and several sensor components. A software infrastructure was built on top of the Robot Operating System to connect the components installed on the robotic vehicle, control them through a single integrated system, and take measurements. Two different navigational modes -- semi-autonomous and autonomous algorithms -- were designed and implemented onto the robot. The core functionality of this system focused on navigation and location of the prototype, as well as decision making depending on the scenarios – including obstacle avoidance. Given the constraints, due to the robot's small size, an environment was built to emulate typical scenarios and used for experimentation to obtain the required performance data. The outcome points to the possibility of replicating this study on a smart wheelchair in real scale, with the aim of reducing costs significantly whilst still addressing safety requirements and ultimately enhancing wheelchair users' independence.

Detection and classification of brain haemorrhage

John Parnis **Supervisor: Prof. Ing. Carl James Debono**

Co-Supervisors: Dr Paul Bezzina and Dr Frank Zarb

Course: B.Sc. IT (Hons.) Computing Science

A brain haemorrhage is a rupture of the blood vessels within the brain, and is very often life-threatening [1]. Brain haemorrhaging is the third-leading cause of mortality across all age groups and is mainly caused by haemorrhagic stroke or trauma [2]. There are various types of blood haemorrhage, including: intracerebral, intraventricular, intraparenchymal, subarachnoid, subdural, and epidural haemorrhages. Diagnosis tools used by medical experts to identify the type of pathology include: computed tomography (CT) or magnetic resonance imaging (MRI) scans, lumbar puncture, or cerebral angiography [3]. This study focuses on the use of CT scans as a method for diagnosis.

The detection and localisation of a brain haemorrhage is highly time-critical, as the longer a case goes undiagnosed, the higher the possibility of a fatality [4]. This study focuses on developing an automated software tool to aid radiographers in classifying the type of bleeding present in a series of CT scan slices, and hence localising it. This is accomplished through the design of a computer-aided diagnosis (CAD) system based on a deep learning algorithm. Specifically, this study explores the use and implements a three-dimensional convolutional neural network (3D CNN). A CNN is a network of layers that reduce an image to its most basic features, making classification easier. The convolutional layer is the layer that translates the image to usable data, scanning small sections of the image and assigning them to different filter classes. The same holds for a 3D CNN, except that in this case the kernels move through three dimensions of data and produce 3D activation maps [5].

When building a CNN structure, it is always best to start with as small a setup as possible and gradually expand, increasing layers and units until the validation error stops improving. Architectural optimisations were also used in this study to improve computational performance. These included the use of skip connections and pointwise filters.

To train and test the developed solution, this study makes use of a dataset of 143 cases, where each case holds multiple slices of CT scan images. In preparation for the next stage, each image was reduced to a 128×128 pixel image and stacked, in order to obtain 143 stacked 3D images with a shape of 100×128×128×1, where 1 represents the greyscale channel and 100 is the stack height. In cases that did not contain precisely 100 images, the available images were truncated or padded accordingly. During this process an image augmentation was also undertaken so as to obtain another 65 images for each training case. The dataset was split into: 80% training and 20% for testing. The divide was carried per classification, rather than as a whole, to ensure a proportionate distribution. The derived results obtained were deemed satisfactory for the purposes of the project.

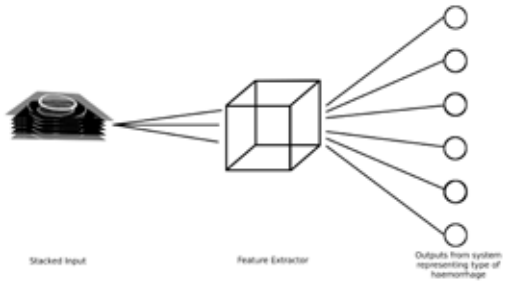

Figure 1. The 3D CNN architecture for the classification of brain haemorrhage from CT scan images

References/Bibliography:

[1] "Cerebral Hemorrhage," accessed 27th August 2020. [Online]. Available: https://www.merriam-webster.com/medical/cerebral%20hemorrhage

[2] N. Spiteri, "Classification of brain haemorrhage in head ct scans using deep learning," Final Year Project, University of Malta, 2019.

[3] A. Felman, "What to know about brain hemorrhage," 2019, accessed 27th August 2020. [Online]. Available: https://www.medicalnewstoday.com/articles/317080

[4] M. Grewal, M. M. Srivastava, P. Kumar, and S. Varadarajan, "RADNET: radiologist level accuracy using deep learning for HEMORRHAGE detection in CT scans," CoRR, vol. abs/1710.04934, 2017. [Online]. Available: http://arxiv.org/abs/1710.04934

[5] "Understanding a 3D CNN and Its Uses," accessed 27th August 2020. [Online]. Available: https://missinglink.ai/guides/convolutional-neural-networks/understanding-3d-cnn-uses/

Enhancing an existing patient dashboard with the use of internet of things

James Theuma | **Supervisor: Dr Conrad Attard**

Course: B.Sc. IT (Hons.) Software Development

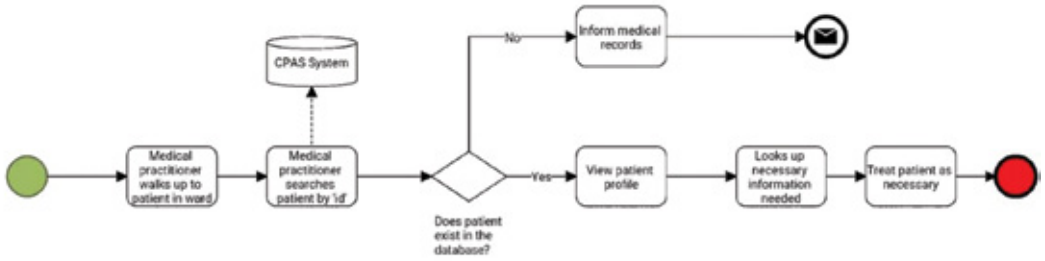

Figure 1. Business-process diagram outlining the ward-round process

The internet of things (IoT) has become increasingly popular within the healthcare sector as it could help resolve various issues and facilitate the automation of a number of tasks [1]. The aim of this dissertation was to enhance the patient dashboard – a web application used by medical staff at Mater Dei Hospital (MDH) – by improving the overall user experience for the said staff, while also applying case studies that make use of IoT.

The patient dashboard connects multiple modules from all departments to store all available medical data regarding individual patient's – such as administered medication, blood test results, etc. – to be displayed on a single screen. Many case studies were viewed and analysed to identify where IoT was being used in the healthcare sector, and how it could be better applied to the patient dashboard [2].

The project also required conducting usability studies. These were carried out among medical practitioners, and also involved developers from the hospital's IT department working on the existing patient dashboard. This was done in order to get a better understanding of how and where IoT is being used at the hospital. Results from this study have shown that, overall, the application offers a wide scope of functionality. However, the medical practitioners also pointed out certain user interface (UI) issues affected the number of steps required to carry out a task. Features such as searching for a patient were only available in the main menu, and navigating through a patient list was not possible. These UI issues proved to be tedious for the medical

Figure 2. Case study example: writing data onto an NFC tag

practitioners when conducting ward rounds at MDH. Upon getting further clarification, the case studies that were given the highest scores by the participants were designed and implemented creating prototypes to be evaluated later on in the study. The new features that were introduced made use of IoT, including QR and barcode scanners, as

well as near-field communication (NFC) protocols.

Following the above, another usability review was carried out on the developed prototypes. The participants found the new UI to be less time-consuming (i.e. in accessing certain features) and also relatively clear and easy to read. Finally, upon testing the features that made use of IoT, medical practitioners were greatly satisfied with the performance and the extent to which the features promised to facilitate their daily task.

References/Bibliography:

[1] S. Madakam, R. Ramaswamy and S. Tripathi, "Internet of Things (IoT): A Literature Review", Journal of Computer and Communications, vol. 03, no. 05, pp. 164-173, 2015. Available: 10.4236/jcc.2015.35021

[2] F. Hu, D. Xieand S. Shen, "On the Application of the Internet of Things in the Field of Medical and Health Care", 2013 IEEE International Conference on Green Computing and Communications and IEEE Internet of Things and IEEE Cyber, Physical and Social Computing, 2013. Available: 10.1109/greencom-ithings-cpscom.2013.38

Data Visualisation using BI for Digital Health

Redent Zammit **Supervisor: Dr Conrad Attard**

Course: B.Sc. IT (Hons.) Computing and Business

The demand for outpatient services is constantly on the increase. Catering for a high volume of patients inevitably causes the said services to encounter difficulties, such as clinic inefficiencies [1]. To mitigate this, the said services have undergone significant changes over the last 20 years [2]. Furthermore, pace of change quickens when both the patients and medical staff make use of the relevant technology to better address the fluctuating rates of referrals from general practitioners (GP) that tends to exhaust specialist resources [1].

This study consists of analysing clinical processes, designing dashboards and developing a solution by applying prediction algorithms. Clinical-appointment scheduling processes have been designed in business-process model notation (BPMN) to better understand how appointments and referrals are performed in the clinic. Business intelligence (BI) dashboards were designed to represent such information in an effective format, and to facilitate clinic analysis. Furthermore, prediction algorithms were applied to the dataset to generate the likelihood of a patient missing an appointment, thus triggering effective future appointment scheduling and helping to prevent missed reschedule opportunities [5, 6, 7].

A case study that has been chosen for the proof-of-concept attempts to address appointment scheduling problems within the Urology Clinic at the Outpatients Department at Mater Dei Hospital. These concern the appointments that are often assigned to non-urgent cases, and are usually not scheduled according to criteria such as patient appointment urgency or age of patient age, among others. Additionally, certain non-urgent, follow-up appointments could be postponed for the benefit of the truly urgent cases. The prediction algorithms enable clinicians to plan for future appointments by knowing which appointments are likely to be missed, and consequently to reschedule the urgent cases accordingly. Moreover, dashboards portraying valuable information would assist the staff in understanding the state of the clinic and patient throughput details [3, 4]. Such patient demographics could aid the consultant and other medical staff to set criteria for incoming appointments resulting from GP referrals.

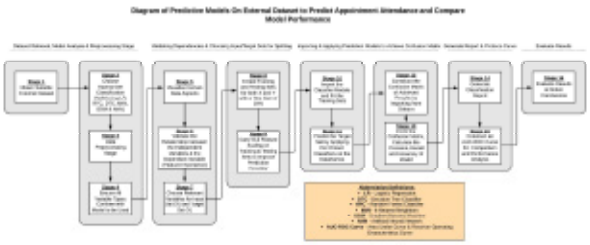

Figure 1. The process of applying prediction models on a publicly available dataset, split into stages

The requisite data for the project, including information regarding requirements at the clinic, were obtained by organising meetings with the consultant and other key medical professionals working at the clinic. An interview was set up with the clinic consultant to discuss in detail clinical processes, current issues and key performance indicators that would best be visualised in dashboards. A usability study was conducted with five participants to gather their feedback regarding the dashboards, their functionality and design. Upon reviewing the available feedback, it was noted that the various respondents were in favour of such dashboards, as they found them to be of significant assistance to medical staff. Since the usability score was very high, it was concluded that the dashboards were suitably designed for this scenario.

At the time of the study, the clinical patient administration system (CPAS) under review, did not capture the required amount of clinical data. Hence, an external dataset had to be sought for the purpose of testing the prediction models. The methodology for applying prediction models on the dataset was divided into stages, as shown in Figure 1. The performance of these prediction models yielded very satisfactory results. The ANN model appeared to be best suited to function on unknown data, and hence could be considered as a possible solution for managing appointment allocation at the Outpatients Department. The outcome of this study could have been more realistic and applicable, had the CPAS captured more data with regard to clinic appointments, schedules and GP referrals.

References/Bibliography:

[1] E. Winpenny, C. Miani, E. Pitchforth, S. King and M. Roland, "Improving the effectiveness and efficiency of outpatient services: a scoping review of interventions at the primary– secondary care interface", Journal of Health Services Research & Policy, vol. 22, no. 1, pp. 53-64, 2016. Available: 10.1177/1355819616648982.

[2] Babiker A, El Husseini ME, Al Nemri A, Al Frayh A, Al Juryyan N, Faki MO, Assiri A, Al Saadi M, Al Zamil F., "Health care professional development: Working as a team to improve patient care." Sudan J Paedi atr 2014; 14(2):9-16. Available: https://www.ncbi.nlm.nih.gov/ pmc/articles/PMC4949805/pdf/sjp-14-9.pdf.

[3] S. Buttigieg, A. Pace and C. Rathert, "Hospital performance dashboards: a literature review", Journal of Health Organization and Management, vol. 31, no. 3, pp. 385-406, 2017. Available: 10.1108/jhom-04-2017-0088.

[4] K. Lee et al., "A novel concept for integrating and delivering health information using a comprehensive digital dashboard: An analysis of healthcare professionals' intention to adopt a new system and the trend of its real usage", International Journal of Medical Informatics, vol. 97, pp. 98-108, 2017. Available: 10.1016/j.ijmedinf.2016.10.001.

[5] D. Hanauer and Y. Huang, "Patient No-Show Predictive Model Development using Multiple Data Sources for an Effective Overbooking Approach", Applied Clinical Informatics, vol. 05, no. 03, pp. 836-860, 2014. Available: 10.4338/aci-2014-04-ra-0026.

[6] C. Elvira, A. Ochoa, J. Gonzálvez and F. Mochón, "Machine-Learning-Based No Show Prediction in Outpatient Visits", International Journal of Interactive Multimedia and Artificial Intelligence, vol. 4, no. 7, p. 29, 2017. Available: 10.9781/ijimai.2017.03.004.

[7] I. Mohammadi, H. Wu, A. Turkcan, T. Toscos and B. Doebbeling, "Data Analytics and Modeling for Appointment No-show in Community Health Centers", Journal of Primary Care & Community Health, vol. 9, pp. 1-5, 2018. Available: 10.1177/2150132718811692.

Transmission of MPEG-DASH over LTE

Max Bonavia **Supervisor: Prof. Ing. Carl James Debono**

Course: B.Sc. (Hons.) Computer Engineering

For this project, a long-term evolution (LTE) communications testbed was used to study and analyse the streaming of videos using MPEG-DASH (Moving Picture Experts Group Dynamic Adaptive Steaming over HTTP) standard. This is an adaptive bitrate streaming technique used by many streaming platforms such as Netflix and YouTube.

The LTE testbed was implemented and configured through the support of Vodafone personnel. The router provided by Vodafone was used to connect a device (in this case a laptop) to the Web through the 4G network. The router itself contains a SIM card to allow access to the network, whilst another SIM card is used to record the received signal strength using a cellular device through the G-NetTrack Lite application. The accompanying diagram illustrates the path of the connections.

This study focuses on the use and benefits of using MPEG-DASH as an adaptive video streaming technique. Hence, any impairments on the audio aspect have not been considered. The different open-access videos used in the project were chosen from those made available by the Institute of Information Technology (ITEC), and were encoded using FFmpeg. FFmpeg is a very convenient open-source tool offering many useful libraries, including libx264 and additional subcommands that specify how the video is to be encoded.

The videos used were from three main categories, namely:

1. Sports – for fast-changing pixels in frames;
2. Film/interview – for a slower-paced but instant frame change;
3. Animation – to analyse a more artificial picture.

Performance curves were plotted to provide a visual assessment of the optimal parameters and limitations of the LTE system being tested. These performance graphs are the result of video frames versus important parameters, such as peak signal-to-noise ratio (PSNR), structural similarity index metric (SSIM) and derivations from the two. Using the open-source video-quality measurement tool (VQMT) executable through command prompt, Excel sheets were created with the parameter levels with every frame of the video, where the original video was compared against the recorded video. This provided a helpful representation of the quality of the transmitted video.

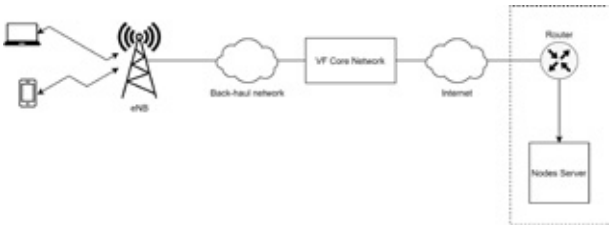

Figure 1. Outline of the system

The objective metrics provided by VQMT are the following:

1. PSNR: peak signal-to-noise ratio
2. SSIM: structural similarity index metric
3. VIFp: Visual Information Fidelity, pixel domain version

VQMT was used on the videos within the following parameters:

1. Perfect conditions: optimal internet speed with a free, unoccupied bandwidth and sound-to-noise ratio (SNR) of 30dB or higher (excellent channel conditions). PSNR range 40-60dB.
2. Bad-mid traffic with good channel conditions: internet speed ranging from 1Mbps to 20Mbps and SNR of 30dB or higher. PSNR range 10-30dB.
3. Good traffic with bad-mid channel conditions: internet speed of 30Mbps and higher, with an average speed of 50Mbps and SNR in the range of 0dB to 20dB. PSNR range 20-45dB.
4. Bad traffic with bad channel conditions (worst conditions): a combination of the bad conditions from Points 2 and 3 in this list. PSNR range: 5-15dB.

For each condition, a range of values using VQMT were obtained. Perfect conditions resulted in PSNRs in the range 40-60dB, whereas the worst conditions only managed 5-15dB. Bad-mid traffic conditions with good channel conditions performed worse than when there is good traffic coupled with bad-mid channel conditions with PSNR values in the range 10-30dB and 10-45dB, respectively. SSIM gave quite similar results in all cases, as envisaged. SSIM should not be below 0.8, and if this occurs it is most likely that the two videos are no longer in phase. Therefore, in such cases the reference frames were compared to wrong output frames.

References/Bibliography:

[1] S. Lederer, "Why YouTube & Netflix use MPEG-DASH in HTML5", Bitmovin, 2015. [Online]. Available: https://bitmovin.com/status-mpeg-dash-today-youtube-netflix-use-html5-beyond/. [Accessed: 18- Aug- 2020].

[2] S. Lederer, C. Müller and C. Timmerer, "Datasets | ITEC – Dynamic Adaptive Streaming over HTTP", Dash.itec.aau.at. [Online]. Available: https://dash.itec.aau.at/dash-dataset/. [Accessed: 17- Aug- 2020].

Transmission of 360 Degree Video over LTE

Damian Debono **Supervisor: Prof. Ing. Carl James Debono**

Co-Supervisor: Dr Mario Cordina

Course: B.Sc. IT (Hons.) Computer Engineering

Advancements in technology have led to new image and video-capturing devices that can record the whole surrounding. Such devices are called 360-degree or omnidirectional video cameras. Moreover, 360-degree video cameras operate at much higher data rate than traditional video cameras, as the data recorded consists of four times more horizontal pixels compared to a typical video. Therefore, effective video-streaming and encoding technology is required to transmit a video stream of such magnitude over a wireless connection.

Most 360-degree cameras currently available on the market transfer the stream of data to a nearby device over a wireless fidelity (Wi-Fi) connection. A problem surrounding this technology occurs when the camera moves further away from the Wi-Fi device, thus weakening the connection and impacting the quality of the received video stream. A solution for such an issue would be to opt for video streaming through a long-term evolution (LTE) connection, as such networks provide better radio coverage that allows for consistent data-transmission rates.

This project involved reviewing a standard approach towards streaming 360-degree live video through LTE technology. The setup consisted of a local cell (eNodeB) that received an ongoing stream of video data from a public server connected to the internet. A set of three different video samples – which differ in bitrate depending on the chosen encoding scheme – was used throughout the research. Moreover, each video sample was transmitted at three different resolutions (FHD, UHD and 4K), thus covering every possible variation that could occur in the videos.

Additionally, a series of both hardware and software techniques were implemented to emulate multiple users and external interferences, as the setup offered a perfect connection unlike real-world scenarios. Hardware techniques included the introduction of attenuation on the signal itself, thus representing signal degradation as the device moved further away from the radio cell. Moreover, dedicated software was utilised to introduce background traffic, as typically the same cell delivers and receives other data from different users, and also to effectively simulate multiple users connecting to the same server. As regards optimal quality of experience (QoE), the system requires that the source upload and user download bandwidths should always be greater or equal to the video-stream bitrate, multiplied by the number of users watching the stream. Furthermore, to determine QoE, the received video was compared with the reference set at the server. This was done through a series of objective measures such as bitrate, peak signal-to-noise ratio, and structural similarity for different setup configurations.

In conclusion, after processing the gathered information, it was found that the statement previously mentioned with regard to system requirements only holds true in perfect network conditions. As traffic and network attenuation were introduced, the stream suffered drastically and QoE deteriorated. This was due to the fact that the network bandwidth was being overloaded with information, which resulted in data packets being lost in transmission. In order to avoid such occurrences, it would be necessary to always leave a portion of the network bandwidth unused, and the signal-to-noise ratio of the received signal should also be above a predefined level, depending on the system being used.

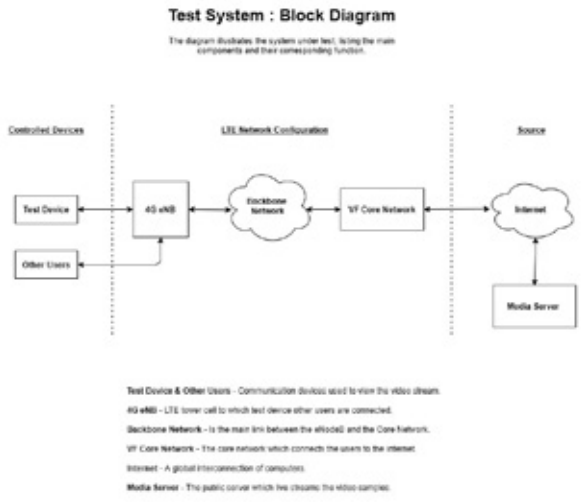

Figure 1. The test System

Interfacing Sensors and Approximating their Position

Carl Joseph Vella **Supervisor: Prof. Ing. Saviour Zammit**

Course: B.Sc. (Hons.) Computer Engineering

The objective of this project is to compare the functionality and performance of three low-power wide-area network (LPWAN) / Internet of things (IoT) technologies – namely, Sigfox [1], LoRa [2] and NB-IoT [3] – which were deployed in a rural area in Malta covering 7 km². The main focus was to conduct experiments, based on three System Development Kits (SDKs) in order to analyse the performance of these technologies in terms of path loss, received signal strength indicator (RSSI) and signal-to-noise ratio (SNR) in outdoor scenarios. In addition, the technical differences between these technologies have been studied theoretically and evaluated in real-world deployments. Finally, the project also considered different application scenarios and elaborated upon which of the three LPWAN technologies would best fit.

Figure 1. Google Earth view of the coverage area

The Hata propagation model was used to estimate path loss for the three technologies. This model predicts the path loss between the end-device and the base station, depending on the height of transmit-and-receive antennae, the frequency of operation, the distance between them, and the type of environment (urban, suburban, rural) and line of sight. The said model was also used to observe how the RSSI correlates with path loss and distance values. The Hata model predicted the values of the path loss consistently for all three technologies. In all the cases, the RSSI was observed from the statistics of their server and compared with the calculations of the Hata model, and the results were found to be reasonably accurate, with quantifiable discrepancies.

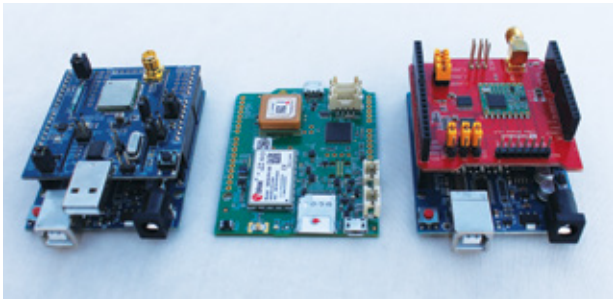

Figure 2. The SDK microcontrollers used in this study

The MATLAB programming language was used consistently throughout the study code scripts to collect data from the three SDKs. The comparison of data was made more statistically accurate by coding scripts to obtain data at certain intervals, and over a long period of time for all the three technologies. This process provided a large sample of compiled data, thus better determining the average values and avoiding errors inherent when testing on a small sample size. Furthermore, MATLAB scripts were alsoused to present statistical data to better understand the behaviour of each individual device in their respective locations.

It is being envisaged that future users could make use of the results obtained from this study to make an informed decision on which technology would be best suited for their preferred IoT application.

References/Bibliography:

[1] J. C. Zuniga and B. Ponsard, "Sigfox system description," LPWAN@ IETF97, Nov.14th, vol. 25, 2016.

[2] U. Noreen, A. Bounceur and L. Clavier, "A study of LoRa low power and wide area network technology," in 2017 International Conference on Advanced Technologies for Signal and Image Processing (ATSIP), 2017, pp. 1-6.

[3] R. Ratasuk et al, "NB-IoT system for M2M communication," in 2016 IEEE Wireless Communications and Networking Conference, 2016, pp. 1-5.

Simulation of Radio Wave Propagation in the CERN Particle Accelerator Complex

Christian Vella **Supervisor: Dr Gianluca Valentino**

Course: B.Sc. (Hons.) Computer Engineering

The aim of this Final Year Project was to study and propose a model to contribute towards optimising wireless communication within the complex of the particle-accelerator tunnels operated by the European Organisation for Nuclear Research, better known as CERN (from the French Conseil européen pour la recherche nucléaire).

This study seeks to follow the development of a model based on a ray-tracing algorithm, to calculate the effects of radio propagation within the field of a tunnel environment. The simulator utilised in the project provided a heat map describing the signal intensity throughout the tunnels given a known transmitter location. Consequently, this facilitated the extracting of the knowledge required to be in a better position to model and plan the best transmitter locations towards optimising coverage. In addition to the aforementioned points, different scenarios were considered, and this was achieved by placing the transmitter at various locations within the context of this project. The overall electromagnetic interference map resulting from the propagation within the network was generated vectorially by combining the resulting propagation maps. This study presents the findings from the different scenarios created during the research process.

The motivation behind the project is to facilitate the work of staff operating within the accelerator-tunnel community, thus increasing the efficiency of their output. In the case of tunnels with considerable length, such as the Large Hardon Collider (which, being approximately 26km is currently the largest particle collider in existence) communication would be crucial in the day-to-day running of the tunnel. Therefore, the model being proposed would contribute towards improving communication between workers possibly located kilometres away from one another. Additionally, the model could also be used to pinpoint any faulty equipment inside a tunnel more effectively

While this project may seem to be bound to one particular tunnel environment, it could be repurposed in various contexts. The project sought to develop a working analytical tool that would allow the replication of wireless propagation in any indoor environment. This would make it possible to easily modify the frequency of the waves, or to implement different materials by changing the dielectric constant.

This project also seeks to ensure that designers and planners of such networks would have all the information they would require in planning such networks. By simulating and testing each set-up, contributes towards eliminating guesswork in choosing different configurations.

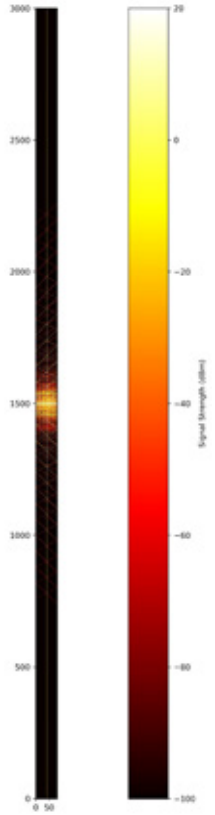

Figure 1. Ray transmitter placed in the centre of the 300m x 8m x 8m tunnel used in the project

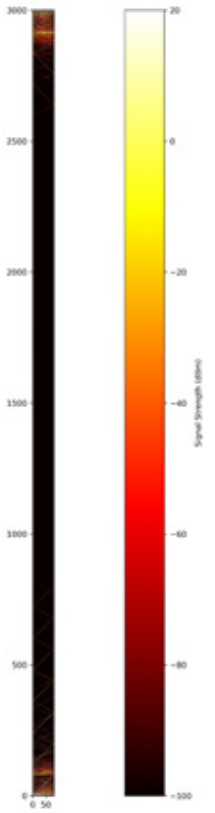

Figure 2. Two ray transmitters placed in two corners of the 300m x 8m x 8m tunnel used in the project

Developing accessibility through computer information systems

Words: Iggy Fenech

ALISON MARIE CAMILLERI and KAREN DIMECH, two students reading for a Master's in Computer Information Systems, are looking to make maps and the internet more accessible. Here's how.

Alison Marie Camilleri

How much do we take for granted? The answer is probably the vast majority of actions we undertake on a daily basis: from watching TV to heading to the nearest corner store, we never bat an eyelid at how many of our senses and abilities we use in order to make that 'simple' action happen. But, for example, how can you use the internet if you are not physically able to manoeuvre a mouse or use a touchscreen? And how do you read a map if you are visually impaired?

These questions have been asked by two Master's in Computer Information Systems students who are currently working on two separate projects that can make the world a whole lot better for those they are aimed at.

The first is by Alison Marie Camilleri, and it seeks to make the internet accessible to a whole host of people who are currently unable to access it due to their physical disabilities. To do this, Alison has created a new type of browser that works using Steady State Visually Evoked Potentials (SSVEPs), which refer to on-screen stimuli in the shape of buttons that flicker at frequencies greater than 5Hz.

"The first part of the study sought to determine whether web technologies could be applied to producing the necessary stimuli for use in a brain-computer interface system," she explains "This included a series of lab-based studies through which we also determined that specialised approximation techniques could allow us to present a greater number of concurrent on-screen stimuli.

"Secondly, having considered a number of factors, such as the stability and the accuracy of generated stimuli, the burden of stimuli-generation on CPU/GPU usage, as well as performance under high-load situations, Cascading Style Sheets (CSS) was identified as the most suitable technology for producing SSVEP stimuli for the web browsing application," she continues.

Currently, the system can determine which of these stimuli the user is focusing on through the processing of brain signal data that is captured via an EEG headset. Through this, users are now able to scroll, refresh, zoom in and out, click on links, interact with input text fields, go backward and forward, type using an on-screen keyboard, and bookmark websites, among other abilities. And all this, without the need to move a mouse, type on a physical keyboard, or tap on a screen.

The browser now supports most standard web browsing functionality, but there is still one major pitfall: the EEG headset is expensive to buy. For that reason, in the future, Alison would like to look into cheaper alternatives for hardware.

The second project is by Karen

Dimech, who is working on making digital maps accessible to visually impaired users. This is being done through a simple yet revolutionary system which gives each element on a map a sound that is triggered when the user's finger is placed on it on a touchscreen. An example of this would be the sound of steps when the user's finger is on a street that has pavements.

The final result is that, through one of these auditory maps, visually impaired users can get from Point A to Point B with greater ease. And this is done not just by the in-app sounds, but by the other things the map gives them.

Users, in fact, can know what type of shops are nearby, thus helping them better orient themselves through real-life sounds and smells. Moreover, they are also told whether there is a pavement on a certain street, and whether there are zebra crossings or steps on their route. They will also be able to know what public amenities are nearby including gardens where they can take their guide dogs.


Karen Dimech

"Karen Dimech [...] is working on making digital maps accessible to visually impaired users"

"At the moment, no popular map service offers this sort of accessibility, which means I had to plot the map from scratch," Karen continues. "Then we had to decide which sounds would work best with each element on the map. We also had to keep in mind that the devices this sort of map would work on couldn't be too expensive.

"The final result is a browser-based map that works well on a variety of devices and which is densely populated with elements that explain exactly what the user will encounter on their journey including any hazards or obstacles like benches, trees, shrubs, steps or bins."

These projects show just how wide the scope of the Information & Communications Technology can be, and that no problem is too abstract to be solved given the right tools are provided and the right amount of determination is present. The best part, however, is that through Alison and Karen's projects, the lives of thousands of people could be made much better in the near future.

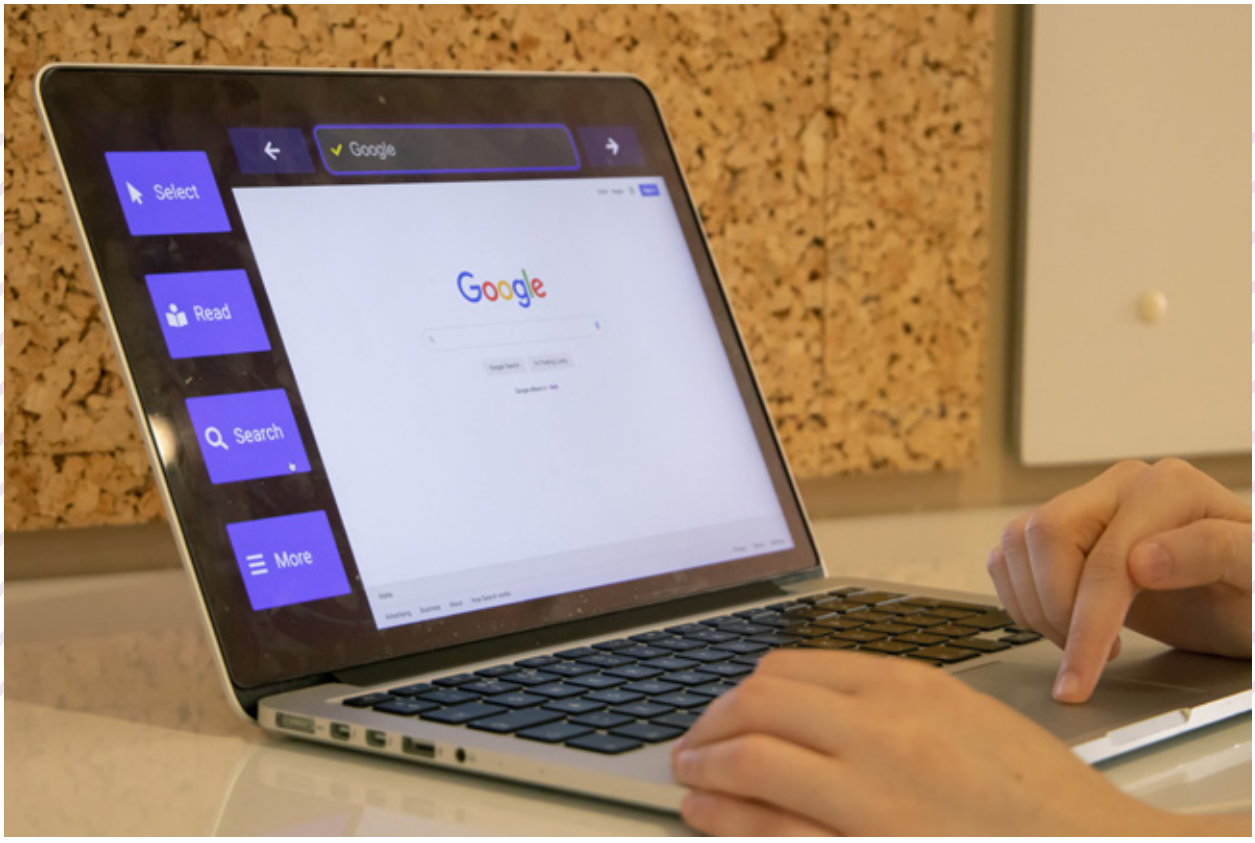

Simulating schizophrenia

Words: Iggy Fenech

By bringing together AI and VR, ANDREW CACHIA is giving a whole generation of student nurses the chance to understand schizophrenia in a way that has never been possible before.

Andrew Cachia

From experts giving advice on various media to celebrities revealing their conditions, mental health is finally being taken seriously and discussed publicly. Yet, while awareness has increased, the truth remains that it is hard for those who do not suffer from mental disorders to understand what a person who does goes through.

That, however, is where Andrew Cachia's Master's in Artificial Intelligence (AI) comes in.

Working together with the Department of Mental Health at the University of Malta, which is the institution that first asked for such a project to be initiated, Andrew has married AI with Virtual Reality (VR) to create a virtual world. In this world people can experience the realities of schizophrenia, a condition in which sufferers interpret the world abnormally, at first hand.

"Once the user enters the simulation, they will find themselves in an office scenario," Andrew explains. "The user is then free to roam the space, interact with the characters, and even use a number of the objects found within it, such as the telephone. The AI side of this comes with the actual interaction."

Upon the user interacting with an object or a character, the AI algorithm, which was created using AI Planning techniques, sets in motion a series of events that are unique to each player's interaction. So, for example, if the user speaks to the colleague character, and the latter asks how the user's day was, the software will initiate a reaction that is correlated to the specific answer.

"So should the user tell the colleague that their day was good, then the user may start hearing bodyless voices repeating how the person's day can never really be good," Andrew continues. "AI truly increases the immersivity of the experience and gives users a unique experience every time they try it."

The idea behind this software is that users get to experience what it's really like to live with schizophrenia. The 'effects' in this virtual world include seeing things that are not there (to correlate with the visual hallucinations schizophrenia sufferers experience) and hearing voices (this applies for auditory hallucinations), among other things.

The main users of this schizophrenia simulation software are student nurses reading for a Master's in treating and handling patients with mental disorders, and it's already being used for this purpose. Many, in fact, have hailed it for the way it humanises the condition and for how it has aided them in being more emphatic in the way they view it.

> “It made people understand the reality behind a mental illness that can be very debilitating”

As a pilot project, this software has truly been a success, and Andrew is incredibly pleased with the results. But, more than just being grateful for what he’s achieved, he’s also appreciative of what the process has taught him.

“I found this project to be extremely interesting,” he says. “Firstly, it combined VR with AI, which is a field that is still largely unexplored. Secondly, there was the challenge of realistically simulating a condition like schizophrenia, which until now has mostly been done through role play. Thirdly, this is contributing to aiding those with mental health issues, which is a noble cause and that makes it all the better in my eyes.”

Talking about what this project could develop into in the future, Andrew says that he would love to continue building on the work that has been done so far. For a start, he would like to create a more realistic virtual world so that the divide between the real and the fictitious is lessened. He’d also like to see the project applied to other scenarios and types of mental disorders so that health students can understand the conditions better, and patients can receive better treatment and more empathy.

“Through the trial runs, we realised how impactful such a project can be. It made people understand the reality behind a mental illness that can be very debilitating.”

Such projects continue to show that learning in the future will be very different from what we experience today but, in the end, it will always serve the same purpose: to make us better, more understanding human beings.

Giving images a voice

Words: Iggy Fenech

Images dominate our world, but most software still can't accurately read them. Through Deep Learning, BRANDON BIRMINGHAM's PhD is set to change that.

The majority of us experience reality in terms of images: what we see shapes the way we interact with the world. Even so, the fact that our brains can easily interpret what we're looking at without much action from our part is something we rarely take into account. But what happens when computers 'look' at images? Well, at this point in time, computers cannot truly understand what an image is of and can't accurately describe it. Nevertheless, as Brandon Birmingham's PhD in Communications & Computer Engineering takes shape, that is all set to change.

Brandon is working on creating Artificial Intelligence (AI) software that can recognise not just what is in an image, but that can also understand the setting, the ambiance and the relationship between the elements within that said image. Just like humans would.

"The idea here isn't just to get the software to say that Picture X is a picture of a cat, but to get the software to interpret the photo in the same way a human being would," Brandon explains. "Where is the cat? What does it look like? Is it asleep, yawning, lounging, eating? And what's surrounding it? Is it all just backdrop or is there anything related to it that is of interest to the viewer?"

One of the main challenges of getting this done is to get the software to see two-dimensional images in a three-dimensional manner, similar to how we, as humans, see

the world. And, surprisingly, one of the hardest things to teach a machine is the difference between prepositions, which it can then use to determine where an object sits in relation to another.

"We are using neural networks to get the machine to learn spatial prepositions," Brandon continues. "These neural networks are basically algorithms inspired by the architecture of the human brain that can be trained to identify underlying relationships between things. In this case, two or more objects in a photo."

Predicting spatial prepositions between image objects, however, is only the beginning as Brandon's plan is to create software that can autonomously create rich and detailed descriptions of images. To begin with, he started with a number of images and scoured the web looking for the most detailed descriptions he could find of them through the use of a mathematical model designed to automatically find relevant captions.

> "The idea is [...] to get the software to interpret a photo in the same way a human being would"

The plan is now to extend this web-retrieval based architecture by constantly accumulating knowledge extracted from the vision and language domains through the use of Deep Learning. In contrast to the majority of the current state-of-the-art supervised captioning models, this proposed self-learning based model will be trained in an unsupervised lifelong learning based approach. This autonomy and independence from humans is extremely important to make Brandon's vision a success.

"I am envisaging three main uses for this software," he explains. "Firstly, through the detailed descriptions, we will be able to retrieve images more accurately off the web or personal image collections – it won't be just about keywords anymore, but about rich explanations. Secondly, visual content will become more accessible to the visually impaired as richer descriptions will bring the images to life. It will also add the possibility to help them understand and navigate in the visual world through the use of Artificial Intelligence-based smart-glasses. Thirdly, this is a further step in the continuation of human-to-robot interaction: in the future, we may be able to simply tell our autonomous car to 'park in front of that green door' and it would be able to understand us in the same way a human would."

Of course, mapping the divide between visuals and language is a challenging task, but it's one that could solve numerous shortcomings and problems. And that is the spirit of ICT in our world, because it's the promise of a solution that makes the work worthwhile.

As for Brandon, he's never let a challenge get in his way. In fact, while he began this project for his Master's degree, he realised straight away that it would be quite a feat to get the software to work, which is why he's persevered and turned into a PhD study.


Brandon Birmingham

Exploring what words mean

Words: Iggy Fenech

Can we humans truly understand what computers actually think when they are presented with textual information? Master's in Artificial Intelligence student BRIAN PACE surely hopes so!

Reading a sentence may seem straightforward to us, but there's a lot that our brains have to take in and process when we do so. Indeed, on top of understanding what the sentence is trying to say, we also have to take into account the context the sentence is in and the feelings the sentence is meant to evoke.

Now imagine when it's a machine that's trying to read a sentence and its job is to extract the sentiment – or the feeling – from the said text. Do you just teach it that the words 'unpleasant' and 'horrible' both infer negative sentiments, even though one is worse than the other? And how would it know that a person had a good experience even though they used the word 'bad' in the sentence, such as in, 'I had previously had a bad experience, but I think the company has come really far and I like it now'?

Well, this is the project Brian Pace is currently working on. And the way he is trying to teach a machine how to decipher the sentiment behind a sentence is through Explainable AI (xAI), which is a set of frameworks that aid those working with them understand and interpret the predictions made by machine learning models.

"While sentiment classification tasks already exist, these are usually based on extracting textual features from the provided data," Brian explains. "This can be achieved through manual processes, such as annotating words in a sentence as content words or high-frequency words, and then using pattern-matching to detect the correct pragmatic import portrayed by the author.

"What I am trying to do, however, is to get xAI to understand for itself whether a sentence is negative, positive or neutral. To do this, it will analyse the robustness of a machine learning model by removing the most prominent features


Brian Pace

from the corpus and observing the decline in accuracy, achieved with interpretability analysis."

To understand this, let's take the sentence, 'I love burgers,' as an example. Through the project, Brian is getting xAI to check how the sentiment of the sentence changes when one specific word is removed. The sentence remains positive both when 'I' and 'burgers' are removed, but changes sentiment when the word 'love' is removed. Therefore, 'love' is the qualifier in this situation.

But, in that case, why not just tell the software that 'love' infers good sentiment?

"One of the problems with current AI systems is that if two words appear frequently together, then the software will start believing that they are both similar qualifiers," Brian continues. "So, in other words, if the words 'love' and 'burgers' appear together often, then the software will classify 'burgers' as a positive word, when, in fact, it's nothing of the sort.

"This doesn't just happen with words, either. In my thesis, I noticed that the software often assumed the hashtag (#) on Twitter is negative, even if the tweet is a positive one. Thankfully, through this system, the software should now understand what it should truly be looking for."

> "I am trying ... to get xAI to understand for itself whether a sentence is negative, positive or neutral"

What we need to keep in mind at this point is that computers don't understand words and sentences the way we do, but they are better with numbers. So the process here is to turn words into numbers. To do so, Brian is using four different types of word embedding software, namely FastText, ELMo, GloVe and Word2Vec. Each of these has its own pros: FastText, for example, takes groups of letters into consideration, while ELMo looks at the text and generates word embeddings on the spot.

This software then works with xAI to give qualifiers a ranking, with 0 being totally negative to 1 being totally positive, thus making 0.5 neutral. In this case, 'bad' isn't as bad as 'horrible', so 'bad' has a ranking of, say, 0.3, while 'horrible' has a ranking of 0.1. Thus, the xAI can then actually tell you whether a sentence is negative, positive, neutral and, more importantly, why it has decided so.

"The benefits of this are numerous," Brian continues. "In the world of words, such software could be used by companies looking to find out what users are saying about their brand or product on social media without the need to have someone doing it manually.

"Nevertheless, such software could also be used for medical purposes by, for example, having the patient inputting symptoms and the xAI accurately telling them what they are suffering from and what medication to take. This, though, requires a lot of precautions and care since the end result will eventually affect the life of a human being, but, in principle, the possibilities are indeed endless."

How such software will impact us in the future remains to be seen, but there's no doubt that xAI coupled with well-thought studies that use its full potential could see the way we live change forever.

Words: Iggy Fenech

Data is king

EMEKA CHUKWU, AKHILESH SHARMA, and LALITMOHAN DIXIT are currently reading degrees in Computer Information Systems. While each of their projects is separate, the one thing that connects them is how they view data.

The 21st century has seen a rise in the value of data. In today's world, data is an abstract currency that can be harvested, analysed, sold and weaponised. Yet data can also be an invaluable resource that helps us gain better insight into a variety of scenarios. At the time of writing, three students within the Faculty of ICT are working with data to ameliorate processes related to healthcare and education.

Emeka Chukwu is the first among these. Currently reading for a PhD in Computer Information Systems, his degree focuses on creating a system whereby health information exchange can happen more seamlessly than ever before.

"The amount of data generated and consumed in healthcare is vast, yet this data is often collected in different formats, either digitally on tablets or computers, or else on paper," he explains. "This means that coordination of data is practically impossible, and it's costing us a lot.

"For a start, it has opened doors to medical fraud and for corporations to make money off healthcare illicitly because no one has a clear, full picture of all that is happening. Worse than that, it's also leading to higher mortality rates because doctors don't have access to each other's notes on patients, and pharmacists may not necessarily know what other medicines a patient may be on when purchasing from them."

"The amount of data generated and consumed in healthcare is vast"

EMEKA CHUKWU

> “I am currently working on a theory that will see machine learning being used to understand the preliminary signs of an epileptic fit”
>
> **AKHILESH SHARMA**

Akhilesh Sharma and Lalit Dixit

To solve this multi-variant healthcare problem, Emeka and his supervisor have designed and published structured architecture that leverages custom Internet of Things (IoT) hardware and standardised healthcare data. These can then be used at points of care or for infectious disease data tracking, and will include the patient’s medical history and lab test results, doctors’ observations and interpretations, medicines the patient is or has been on, and the likes.

Of course, Emeka understands that traditional trusted data centralisation could be a gold mine for corrupt entities, which is why he is adamant that the model’s final version would have to run on decentralised architecture for shared security and trust. He insists it can only scale with adequate government controls like those in place in Malta and championed by Malta Digital Innovation Authority (MDIA).

“The concept of standardised data exchange isn’t a farfetched idea, either,” he explains. “Banks use standardised safety measures to transfer data and money with great success, so why shouldn’t healthcare? Plus, with so many benefits to people’s health, it’s something worth investing in.”

The idea that data could save lives is a widespread concept. In fact, Master’s in Computer Information Technology student Akhilesh Sharma is using electroencephalogram (EEG) data to understand what happens in the brain right before an epileptic seizure, specifically concentrating on the low frequency brain waves.

“Together with my supervisors, I am currently working on a theory that will see machine learning being used to understand the preliminary signs of an epileptic fit; the changes in the brain that happen before the rolling of the eyes, shaking of hands, and difficulty in speaking associated with such a seizure,” Akhilesh explains.

“This could have multiple benefits. On the one hand, it means that we could better understand how epilepsy affects the brain while, on the other, it could see health workers anticipate such seizures and give them precious time to moderate the effects and even save the person’s life.”

At the moment, Akhilesh’s Master’s is still at the literature review state, and the actual research has had to be delayed due to COVID-19. Yet this change in the way students all across the world have had to grapple with could itself be a goldmine in data terms, and that is something Lalitmohan Dixit, who is also reading for a Master’s in Computer Information Technology, is keen to tap into.

“My Master’s is looking at neurological learning patterns to understand the actual differences in the ways students learn inside a classroom and at home through online lessons,” he explains.

Also using an EEG, the data collected looks to answer a number of questions including which scenario is best to help students focus and how their brains react to the different stimuli.

“As a Master’s student, I’ve only had one full month of classes so far due to the academic year being disrupted by the pandemic,” he continues. “This, however, has made me understand that there may be a huge difference in how our brains learn when we change scenarios, and I’m hoping this study will help us discover what the best way of learning in the modern world is.”

In all three projects, data, its collection, and its distribution are the key factors, but the results they hope to achieve make these projects very human in nature. Data, after all, is not just about numbers, but also about understanding a segment of people and their needs.

An artificially intelligent helper

Words: Iggy Fenech

Banks are about to get a helping hand when it comes to deciding whether an applicant should be granted a loan or not thanks to LARA MARIE DEMAJO's Master's in Artificial Intelligence.

Any bank will tell you that one of the loan officer's toughest jobs is deciding whether or not an applicant will be able to repay their loan as determined by their contract. With so many variables to look at, it's not hard to see why it sometimes takes ages for a decision to be made. Yet a new credit scoring model currently being worked on by one of the students reading for a Master's in Artificial Intelligence, Lara Marie Demajo, could very well halve the loan officer's job, leaving them with more time to spend on what really matters.

"Using a machine learning model that is equipped with an XGBoost algorithm, I have created a credit scoring system that not only analyses the data inputted into it to come up with a confirmation or denial of an applicant's loan, but which also gives a detailed explanation of the logic it used in making its decision," Lara explains. "This, in my opinion, as well as in that of the bank managers I have spoken to, will result in a more seamless and fair way for technology to help loan officers in their job while cutting waiting time for applicants."

For those who are not familiar with credit scoring, this is a system which is used by banks to decide whether a person should be given a loan or not, as well as how big that loan should be, judging by past credit, the person's wage, and other loans and assets,

among other factors. The current situation, however, sees loan officers often doing all the work themselves, even for cases where the results should be a black-and-white affair.

"The idea here isn't to replace the loan officer but to offer them assistance that can save them precious time on applications which are easy to determine based on the information provided. This means that the loan officers could then use their time to go over applications which are more sensitive or less straightforward."

The explanation given by the software will also prove incredibly useful when speaking to the prospective applicants as it breaks down the reasoning behind its decision. Together, the loan officers and applicants could then discuss what changes would need to be made in order for them to get a loan.

"The idea here isn't to replace the loan officer but to offer them assistance that can save them precious time"

"Of course, the software does not have the final say. Instead, it goes through the information and gives the predicted outcome based on the logic it has learned from past loan applications, which can then be further analysed by the loan officer," Lara continues. "The explanation, meanwhile, also aids banks in observing the laws related to GDPR and ECOA, which require them to share the reason behind their decisions."

For this system to exist, Lara trained the model using large data sets containing information of different loan applications from two foreign credit institutions. Moreover, she analysed if the provided explanations are easy to understand by gathering opinionative results through interviews with seven loan officers, as well as a hundred questionnaires filled in by the general public. As a result, this state-of-the-art system is much more sophisticated than the counterparts currently used by the majority of the banks.

"Needless to say, there is room for improvement in terms of accuracy, but there may be a time when such systems could also help determine the cases which are not so straightforward. Even so, one of the most important things that have come out of my research is the fact that younger generations are more inclined than ever to trust such software with important decisions. Indeed, it shows a clear shift in mindset and it paves the way for the world we'll be living in in the years to come."

In the future, Lara is adamant that more AI systems will use similar approaches not just to make choices but also to explain them. This, for example, would be particularly useful when it comes to autonomous cars, which make decisions intuitively and autonomously that affect both those in it, as well as those driving nearby, pedestrians and so forth.

"Knowing why the AI software decided that Option A was more viable than Option B helps us build a better rapport with it, as well as understand any pitfalls," Lara adds.

Of course, not all situations can be tackled using logic on its own. How would the world fare if medical choices were based solely on logic rather than logic coupled with empathy and humanity? Yet, that doesn't change the fact that having an independent arbitrator can help us make more accurate choices in a shorter time span.

Underneath the bandages

Words: Iggy Fenech

Through the use of microtomography, DR MARC TANTI is looking to bring the study of ancient mummies firmly into the 21st century.

While archaeology seeks to uncover our past, it has to be firmly rooted in the present, particularly as new techniques can help us understand our discoveries better. Even so, just because a technological process aids us, doesn't mean that it makes things quicker. In fact, as things stand, when researchers create volumetric images of mummies (3D renditions of a physical objects created through X-rays), they can only find out what's in each layer by manually segmenting each slice – a process that takes months.

Now, however, thanks to Dr Marc Tanti's contribution for his post-Doc, which saw him work on an international project called Automated Segmentation of Microtomography Imaging (ASEMI), the timeframes of working on such artefacts may indeed be halved.

Together with a team of five academics from the University of Malta, as well as four researchers from the European Synchrotron Radiation Facility (ESRF) in France, Marc has used machine learning techniques to create software that can, almost automatically, determine the type of materials found inside objects, specifically animal mummies. This is done by reading the density and texture visible in the volumetric image that results from microtomography.

"We start off with the actual scanning," Marc explains," which which is done at the ESRF. The process here is to accelerate electrons in a circular accelerator that emits high-energy X-rays, which are then passed through the specimen and onto an imaging sensor. These projectional radiographs are then used to compute a 3D image of the animal mummies, which is called a computer microtomography image or a volumetric image."

Now, you may think that the scanning of mummies and other artefacts has become commonplace across the world of archaeology and beyond, and you wouldn't be wrong. Similar, but less powerful, systems have, after all, been used in material science and healthcare, but it's what will happen with the volumetric image of the mummies that is the real breakthrough in all this.

See, currently, the computed microtomography images of a whole mummy have to be manually seg-

Dr Marc Tanti

mented. This laborious process is particularly time consuming and can take months for some specimens. With the ASEMI technology, however, archaeologists will be able to manually segment just a small number of slices and use the information to teach the software what it is it should be looking for as it segments the rest of the volumetric image by itself.

"In other words, we can know whether there are bones, textiles, metals, hair, nails or even food in the stomach without the need to segment each layer manually.

"This, of course, also saves researchers a lot of time, but higher speed doesn't just mean being able to take more coffee breaks," Marc continues. "It also means that archaeologists and historians can start compiling statistics that would have otherwise been impossible to put together so quickly before: the more mummies are scanned, the better we can understand the practice of mummification."

And mummies have a lot to teach us about the past. One particular mummy of an ancient Egyptian crocodile, whose computer microtomography image took the team at ESRF three months to manually segment, revealed many of its secrets including how it was killed, where it lived, and its diet.

When it comes to ASEMI, Marc's job was to try out a number of options for the process and to report on the results for each to the rest of the team. It must be pointed out, however, that prior to taking up this project for his post-Doc, Marc had absolutely no archaeological knowledge. In fact, he, as part of a team from Malta, visited ESRF to receive training on how to manually segment mummies in order to be able to teach the machine learning software how to do it properly for itself.

"It also means that archaeologists and historians can start compiling statistics that would have otherwise been impossible to put together so quickly before"

Even so, the results of the project speak for themselves and show how multi-disciplinary projects such as these – this particular one brought together physics, archaeology and computing – could help various departments advance their learning and work.

The ASEMI project is funded by the ATTRACT project, which in turn is funded by the EC under Grant Agreement 777222.

Giving the elderly more independence

Words: Iggy Fenech

Autonomous guided vehicles have long been used in factories and warehouses, but could they soon find their way into care homes? Master's student MARK MIZZI is certainly working towards it.

Care homes are a practical and emphatic solution to ensuring that elderly people who can no longer take care of themselves find the help they need to go about their day-to-day lives. Even so, it doesn't change the reality that a number of those who enter care homes may not like the fact that their autotomy is compromised by their age and lack of physical ability. Moreover, while help may be readily given, it often means that the elderly have to wait their turn to be seen to by carers whose tasks and responsibilities are numerous.

Of course, not all of the carers' tasks can be replaced by technology, nor should they be. The human element in these sorts of situations is incredibly important. Yet technology can offer a much-needed helping hand to carers while giving the patient increased independence, particularly when it comes to mobility.

While studying the best way to do this for his Bachelor's degree, Mark Mizzi started looking at how certain industries use autonomous guided vehicle (AGV) systems to facilitate their processes.

"In many factories and warehouses around the world, materials and products are transported around using AGVs that work without the need of human assistance," says Mark, who is now continuing his studies through a Master's in Computer Information Systems.

> "When it comes to navigating a space, [smart] wheelchairs fall short of what they promise"

"These AGVs are usually connected to one of three systems. The first is the centralised approach, where a server has a digital map of the floor and can send out instructions to the AGV with what needs to be collected, where it needs to be delivered, and the best route to take to get the objects from Point A to Point B along with which obstacles will need to be avoided. The second is the decentralised approach, where the AGVs communicate among themselves to decide which one of them is closest to the objects needed and to find the best route to deliver it."

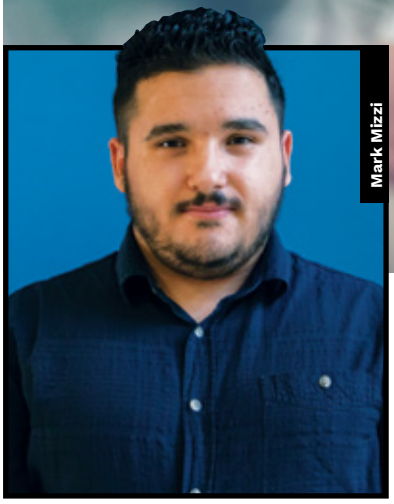

Mark Mizzi

While each of these systems is used by numerous businesses, there are pitfalls. The centralised approach, for example, can have 'system blackholes,' which means the system won't be able to reach certain parts of the premises. The decentralised approach, meanwhile, sees the AGVs solely communicating with other AGVs with no central system overseeing the process, thus resulting in wasted time and a higher number of collisions.

"This is where the third system, the hierarchical approach, comes in," Mark continues. "This approach removes the majority of the cons of the first two approaches by having a centralised system that sends out the

information and maps out the route, and then a decentralised system which sees the AGVs communicating to avoid collision, obstacles, bottlenecks and so on."

You may be wondering what this has to do with care homes, but given that in some care homes up to 50 per cent of elderly residents need to make use of a wheelchair, utilising such systems could have a huge impact on both carers and residents

"Smart wheelchairs have long been researched and used in care homes, but when it comes to navigating a space, these wheelchairs fall short of what they promise," Mark continues. "One of their major shortcomings is that they cannot predict the best route based on traffic, meaning that having a number of them in one care home could see the corridors and halls turn into absolute mayhem that can only be solved by human beings.

"Therefore, one of the main advantages of using the hierarchical approach in care homes would be that residents who need to move from one place to another using a wheelchair would be able to do so by themselves. All without the worry of getting stuck in internal traffic jams or coming across obstacles that the smart wheelchairs can't predict."

These traffic jams are not something that may occur only in the future, either. Through his research, for which he interviewed a

number of managers of local care homes, Mark has discovered that even traditional wheelchairs tend to end up in bottlenecks, often causing loss of time for the carers and an inconvenience to the residents.

Such a system, therefore, would see the residents travel more freely inside the care home, see one of the carers' most time-consuming tasks being removed from their to-do list, and make managing the care home that much easier thanks to the server that oversees the process of residents moving about.

At the time of writing, the hierarchical system Mark is suggesting for care homes is being tested using simulation, but he expects it to be ready for real-life testing in the years to come – and, who knows, it may just give the elderly residents of Malta and Gozo's care homes more independence than is currently possible for them to have.

Safeguarding independence through AI

Words: Iggy Fenech

Artificial intelligence could soon be helping carers take care of the elderly. MATTHEW SACCO, who recently completed his Master's in Computer Science, explains his role in the project.

On average, people live longer now than they did in the past. That is both a fact and a blessing. Even so, an aging population comes with a number of issues, one of which is how we can reach a compromise between ensuring that the elderly are well looked after without taking their independence away.

This is something Matthew Sacco had to grapple with as he joined a team of individuals and entities located in Malta and Sicily to work on an EU project called NATIF Life, whose aim is to create the assisted-living apartment of the future.

"The project trickled down to me as a student doing my Master's Degree in Computer Science," Matthew explains. "From my end, I used a special type of camera termed as RGB-D, which has depth-sensors that can determine how far an object is from it. A system made up of a number of such cameras, coupled with Artificial Intelligence smart enough to detect what the people in a room are doing, means that we can now have an indoor localisation system that can tell when a person is in trouble or is simply going about their day-to-day business."

Setting the camera system up in a laboratory within the Faculty of ICT, as well as at a care home in Sicily, Matthew taught the software how to tell whether a person had fallen down and required help, among other skills. Moreover, by speaking to relatives of

elderly persons who live at Sir Paul Boffa Hospital, Matthew realised that the software could help relatives with some often forgotten issues.

"Many relatives mentioned that some of the hardest things to keep track of are whether the elderly in their care had been active, had eaten, had taken their pills, or done their exercises. It's not something we had really taken into account when we set out to work on this project, but they're pertinent issues, particularly as these all help elderly people remain strong," he continues.

One of the biggest hurdles of the project was to ensure that the privacy of the individuals living in the assisted-living apartments of the future would be safeguarded – something which is already a major headache with CCTV and surveillance cameras.

To achieve this, the first step was to use the type of cameras they did, which do not 'see' colour and don't have the ability to recognise one person from another. Moreover, while the information currently being retrieved from the laboratory here in Malta and the care home in Sicily is going into a centralised system for data collection, this would only be done with explicit consent from the elderly people in question.

Nevertheless, the potential for such a project is huge. Firstly, this could help keep track of how elderly people are doing when no one can physically be around, a particularly important factor for those families which may not afford a live-in carer. Secondly, the system can be programmed to automatically inform the relevant people and authorities should something go wrong – in case of a fall, a fire or a burglary, for example. Thirdly, when used at care homes, the system could act as a second pair of eyes to help carers with their work.

> "In certain situations, such as near staircases, where most accidents tend to happen, this could truly become a life-saving project"

Matthew Sacco

"Indeed, while the system is ready for day-to-day usage, the probability is that it will have its roll out in care homes. There it could truly benefit the elderly as many such homes are finding it hard to recruit enough personnel and it's not always possible to be with each patient at all times. In certain situations, such as near staircases, where most accidents tend to happen, this could truly become a life-saving project."

The system, in fact, is already being used for such purposes at the aforementioned assisted-living property in Sicily, which is a temporary retreat for the elderly who may need a half-way house while recovering, among other reasons.

The best part of the project, however, remains the way Matthew designed the system. Its modular nature means that, in the future, people could create new applications that would integrate seamlessly, thus adding to the range of features the system can support.

"There are already more research studies being proposed in order to update the system and I look forward to seeing what I and other researchers manage to do with this framework," Matthew concludes.

That makes two of us, Matthew!

Environmental sensors for air quality (ESAIRQ)

Words: Iggy Fenech

How good is the air quality inside buildings? DR RUSSELL FARRUGIA, and PhD students BARNABY PORTELLI and MATTHEW MELI, are part of an international consortium looking to create a system to find out.

Air quality is a phrase that has got a lot of traction as of late and as we continue to learn about the health hazards of humans breathing bad air, we can expect that to intensify. Yet it's good to remember that the term 'air quality' doesn't just refer to the outside environment, but also to that inside buildings.

In order to get to grips with how best to monitor and manage the air quality found indoors through gas sensing technologies, a consortium made up of twenty-six partners hailing from six EU member states is now in place. The partners include companies like Philips and Infineon, as well as other research institutes and Universities like Fraunhofer ENAS, the Eindhoven University of Technology, and the University of Malta (UM).

"UM has two main roles in this multinational project entitled ES-AIRQ," explains Dr Russell Farrugia, who has been on the project since its inception. "Firstly, there is the development of a network made up of over 100 air quality monitoring wireless sensor nodes. Secondly, there is the development of a gas sensor based on infrared spectrometry."

For this to happen, the Faculty's building at the University of Malta will now have over 100 such prototype nodes installed with the aim of monitoring environmental values of things like temperature and humidity, as well as that of numerous gases including volatile organic compounds (VOCs), carbon dioxide (CO2) and particulate matter (PM).

"The wireless communication protocol selected for these nodes is LoRa, a Low Power Wide Area Network (LPWAN) protocol developed specifically for the Internet of Things (IoT) to enable low power communication over distances of up to 2km in urban

environments," adds Matthew Meli, whose PhD revolves around the project. "Moreover, given the star topology used, if one sensor node fails, all the other nodes can continue working normally."


Matthew Meli

As one can imagine, with each of the 100-plus nodes collecting data on a regular basis, the final amount of information will be huge. Yet this data is exactly what the team working on the ESAIRQ project is after.

"The big data attained from the sensors is being stored in a cloud-based database," says Barnaby Portelli, who is also reading for his PhD. "This way, we not only have the possibility of making big data analysis, but it also means that our partners on the project can see and access the results of the study. Moreover, our sensor nodes have multiple interfaces available to test additional, newly developed sensors provided by our partners in the project."

Even so, there is one pitfall that the consortium is looking to have sorted out in the near future, which is that of making the battery powered sensor nodes cheaper to buy and longer lasting.

"Part of my research area is to drastically extend the battery life of the sensor nodes by focusing on ultra-low power design," Matthew continues. "At the moment, the air quality sensor nodes that are on the market are not only expensive but also have an incredibly short battery life. Our aim now is to achieve a system where sensor nodes are able to perform air quality measurements autonomously for multiple years."

On top of all this, the consortium is developing an innovative portable device capable of sensing a range of gases using near-infrared spectrometry to detect multiple gases.

"The fundamental component of the gas sensor is a 2mm-diameter scanning diffraction grating designed in-house and fabricated using silicon semiconductor technology," Barnaby explains. "The gases we are looking to detect and monitor include Carbon Dioxide (CO_2) and Methane (CH_4)."

But, you may ask yourself, what are the advantages of such studies?

"Well, through such a project, Europe will gain essential knowledge on gas-, fine-particle-, and pathogen-sensing technologies that could have huge benefits for people's health," Russell continues.

Indeed, this could help spell the end of the sick building syndrome, which is a phenomenon that sees occupants of certain residential or office buildings experience numerous, non-specific symptoms, such as headaches, fatigue, or eye, nose and throat infections. As the name states, this is normally attributed to the bad air quality inside buildings.


Dr Russell Farrugia

Moreover, while speaking to the trio, they all agree that this could be the first step towards having smart, artificially intelligent systems that don't just monitor air quality, but also manage it. In other words, the system itself could identify what harmful gases may be in the air and be capable of remedying the situation by changing the air flow automatically.

> "This could be the first step towards having smart, artificially intelligent systems that don't just monitor air quality, but also manage it"


Barnaby Portelli

Of course, that is something that is much easier said than done, particularly as such a system would also need to tailor its response based on multiple factors including how many people are in the building at any given time.

Even so, the potential of such a system is huge and the benefits to both individuals and business could be even bigger: we could decrease illnesses, increase productivity and brighten people's moods all through a smart system like this. And if that does happen, part of the thanks should definitely go to the Department of Microelectronics and Nanoelectronics within the Faculty of ICT!

What we do

Foster independent, consistent high-quality regulation that drives investment and cultivates innovation in modern communications.

Promote competition and ensure communications markets work effectively for consumers.

Boost eCommerce and digital services growth through an apt regulatory environment.

Ensure availability and sustainability of high-quality universal services fit for today's age.

Safeguard consumer rights and empower them to make informed decisions.

Ensure optimal use of scarce resources, including numbering ranges and radio spectrum, that enable high-tech business and services to flourish.

Support national and EU digital policies through forward-thinking specialised advice for an inclusive digital society.

www.mca.org.mt

MALTA COMMUNICATIONS AUTHORITY

Change your job into a career!

The Faculty of Information & Communication Technology (ICT) offers a range of specialist courses that enable students to study in advanced areas of expertise, and improve their strategic skills to achieve career progression.

Get to know more about our courses
um.edu.mt/ict

- Master of Science in Artificial Intelligence *(Taught and Research, mainly by Research)*
- Master of Science in Computer Information Systems *(Taught and Research, mainly by Research)*
- Master of Science in Computer Science *(Taught and Research, mainly by Research)*
- Master of Science in Computer Science *(Taught and Research, mainly Taught)*
- Master of Science in Signal Processing and Machine Learning *(Taught and Research, mainly Taught)*
- Master of Science in Microelectronics and Microsystems *(Taught and Research, mainly Taught)*
- Master of Science in Telecommunications Engineering *(Taught and Research, mainly Taught)*
- Master of Science in Information and Communication Technology (Computer Information Systems)
- Master of Science in Information and Communication Technology (Microelectronics and Microsystems)
- Master of Science in Information and Communication Technology (Signal Processing and Machine Learning)
- Master of Science in Information and Communication Technology (Telecommunications)
- Master of Science *(by Research)*
- Master of Science in Human Language Science and Technology

RS2

together beyond payments

Ever thought about the technologies making payments convenient, seamless, and secure?

That is our job - We shape the future of payments!

We are

- A leading provider of global omnichannel payment technologies and solutions for acquirers and issuers on one single platform
- Located in 6 offices in Malta, Germany, US, Philippines, Jordan, Brazil
- 380 Employees

We process

- 80 Million transactions per hour
- 350 Million cards
- 200 payment methods
- On a single platform for all payment methods across the globe

We serve

- World's leading Banks
- Financial Institutions
- Payment Service Providers
- Payment Facilitators
- Independent Sales Organizations
- Independent Software Vendors
- 16 ,Million Merchants- indirectly
- in 35 Countries

RS2.COM

CCBill

SKYROCKET YOUR ICT CAREER

employment.ccbill.com

scan me

'PERKS' PACKAGE

- </> Flexible Working Hours
- </> Gym Subsidy
- </> Monthly Technical Workshops
- </> Free Daily Gourmet Buffet Lunches
- </> Online Training Courses
- </> Private Hospital Insurance
- </> Conferences Abroad
- </> Snack Vending Machine Allowance
- </> Free Soft Drinks, Juices, Fruit, Tea & Coffee
- </> Discounted Rates On Laundry

OUR TOOLBOX

CI CD, Angular, NGINX, Bitbucket, GitLab, LINUX, git, spring by Pivotal, ANSIBLE, JIRA, Jenkins, vmware, CSS3, HTML5, ES6, Java, kubernetes, Maven, docker

Interested? Email us: **cindya@ccbill.com**

Parallels

Parallels:
Simplicity over complexity

Parallels Inc., a global leader in cross-platform solutions, makes it simple for customers to use and access applications and files they need on any device or operating system.

Parallels has been a “Gold partner” of the University of Malta since 2016

For the UM students there is the Parallels “Students Partnership Experience” initiative allowing talented young individuals to take part in Work exposure and Internship programs to try themselves with real life projects as part of the RAS team.

Working/Studying from home? Do you need full access to your applications wherever you are?

We make it happen! We have already helped thousands of organizations to easily continue to work remotely. Joins us to help in saving the economy!

Join our team and make it happen!
parallels.com/eu/about/careers

The heart of the Parallels Remote Application Server (RAS) development is in Malta

Parallels RAS is a streamlined remote working solution providing secure access to virtual desktops and applications:

- Deliver virtual desktops and apps to any device, anywhere, anytime.
- Enhance data security by centrally monitoring and restricting access.
- Quickly scale your IT infrastructure on demand with auto-provisioning.

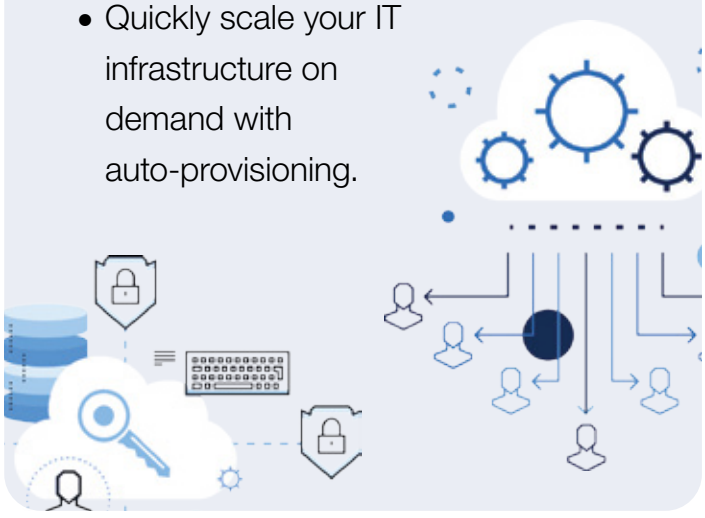

Parallels is an integral part of Corel Corporation.

KPMG

Accelerate your Digital Transformation

KPMG brings an industry-led, customer-centric approach to transformation that aligns the front, middle and back offices for efficiency, agility and sustainable growth.

With technology solutions, frameworks, accelerators and tools designed and engineered to support continuous innovation, sustainable high performance and success in a fast-moving digital world – we help our clients **accelerate their digital transformation.**

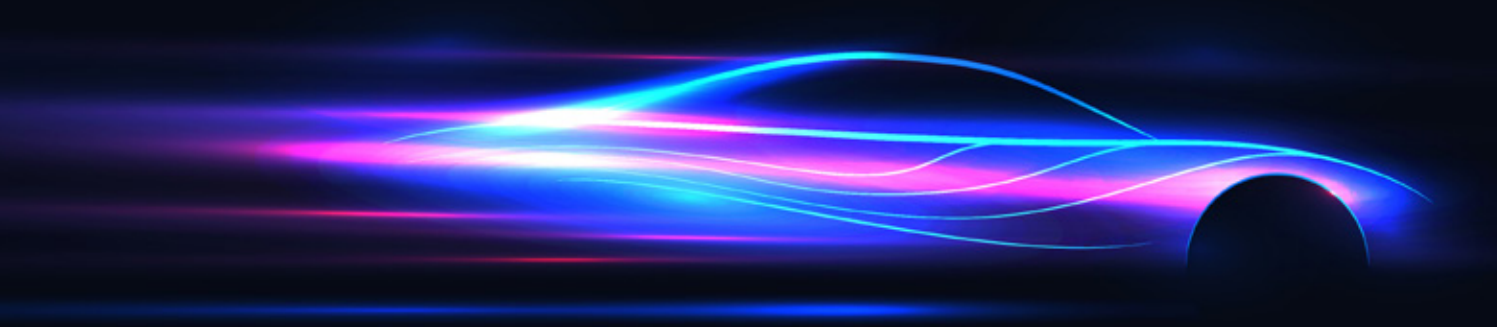

Marco Vassallo
Partner, Digital Solutions
+356 7942 9428
marcovassallo@kpmg.com.mt

Right now, we're all asking ourselves big, complex questions.

Let's work on the answers together.

www.kpmg.com.mt

Artificial Intelligence

ICT

Blockchain

LAUNCH YOUR STARTUP

Digital Games

Software Development

Take your studies and ideas to the next level.

Speak to us to learn how we can help you launch your own successful enterprise.

W: www.maltaenterprise.com
E: info@maltaenterprise.com
T: (+356) 2542 0000

MALTAENTERPRISE

mita

OUR DIGITAL FUTURE

> mita.gov.mt