BABEŞ-BOLYAI UNIVERSITY
FACULTY OF MATHEMATICS AND COMPUTER SCIENCE

Machine Learning and Advanced Computer Vision Techniques for Natural Sciences

Summary of the Ph.D Thesis

Ph.D. student: Horea-Bogdan Mureşan
Scientific supervisor: Prof. Dr. Horia F. Pop

2023

Keywords: Convolutional neural networks, deep learning, synthetic data, object classification and detection from images

Contents

Contents

List of Publications

All rankings are listed according to the classification of journals[1] and conferences[2] in Computer Science from the year of the publication. The evaluation is in accordance to the standard of the CNATDCU committee from 1^{st} of October 2018.

Publications in ISI Web of Knowledge

Publications in ISI Conference Proceedings Citation Index

1. [Mur19] **Horea-Bogdan Mureşan**, Skin lesion diagnosis using deep learning, *IEEE 15th International Conference on Intelligent Computer Communication and Processing (ICCP)*, pp. 499–506, 2019, https://doi.org/10.1109/ICCP48234.2019.8959661.
   CORE 2018, Rank C, 2 points, 7 citations

2. [MCC20b] **Horea-Bogdan Mureşan**, Adriana Mihaela Coroiu, Alina Delia Călin. Detecting leaf plant diseases using deep learning: A review, *2020 International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, pp. 1–6, 2020, https://doi.org/10.23919/SoftCOM50211.2020.9238318.
   CORE 2020, Rank D, 1 point, 0 citations

3. [MMC+21] Alexandru-Ion Marinescu, **Horea-Bogdan Mureşan**, Alina-Delia Căalin, Adriana-Mihaela Coroiu, Maria Talla. Freida - fracture risk evaluation using highly efficient information retrieval and analysis of large healthcare datasets. *IEEE International Conference on Big Data (Big Data)*, pp. 3917–3922, 2021, https://doi.org/10.1109/BigData52589.2021.9671799.
   CORE 2021, Rank B, 1.33 points, 1 citation

[1] https://uefiscdi.gov.ro/scientometrie-reviste
[2] http://portal.core.edu.au/conf-ranks/

4. [Mur22] **Horea-Bogdan Muresşan**. An automated algorithm for fruit image dataset building. *17th Conference on Computer Science and Intelligence Systems (FedCSIS)*, pp. 103–107, 2022, https://dx.doi.org/10.15439/2022F58.
CORE 2021, Rank B, Short paper, 2.66 points, 1 citation
(Reclassified as Multi-conference in November 2022)

5. [CCM22] Adriana Mihaela Coroiu, Alina Delia Călin, **Horea-Bogdan Mureşan**. Collaborative transdisciplinary educational approaches in AI. *International Conference on Computer Supported Education*, pp. 260-267, 2022, https://doi.org/10.5220/0011039500003182.
CORE 2021, Rank B, 4 points, 0 citations

Publications in International Journals and Conferences

1. [MO18] **Horea Mureşan**, Mihai Oltean. Fruit recognition from images using deep learning, *Acta Universitatis Sapientiae, Informatica*, volume 10(1), pp. 26 – 42, 2018, https://doi.org/10.2478/ausi-2018-0002.
UEFISCDI 2018, Rank D, 0 points, 129 citations
note: the preprint of the article available at https://arxiv.org/abs/1712.00580 has 384 citations

2. [MCC20a] **Horea-Bogdan Mureşan**, Alina Delia Călin, Adriana Mihaela Coroiu. Overview of recent deep learning methods applied in fruit counting for yield estimation. *Studia Universitatis Babeş-Bolyai Informatica*, volume 65(2), pp. 50–65, 2020, https://doi.org/10.24193/subbi.2020.2.04.
UEFISCDI 2020, Rank D, 1 point, 0 citations

3. [CCM23] Alina Delia Călin, Adriana Mihaela Coroiu, **Horea-Bogdan Mureşan**. Analysis of Preprocessing Techniques for Missing Data in the Prediction of Sunflower Yield in Response to the Effects of Climate Change. *Applied Sciences*, 2023, 13, 7415, https://doi.org/10.3390/app13137415.
4 points, 1 citation
UEFISCSI 2022 AIS, Rank B (Q3 5-th position, Eng. Multidisciplinary SCIE)
UEFISCSI 2023 AIS, Rank B (Q3 7-th position, Chem. Multidisciplinary SCIE)

Chapter 1

Introduction

This chapter presents the motivation behind the problems approached in this Ph.D. thesis, the objectives that we aimed to accomplish

1.1 Motivation

Computer science has always impacted other research domains, through introducing innovative methods of data processing, and through allowing better and faster ways of understanding data. Moreover, computer science has enabled the automation of tasks that conventionally are time consuming when performed by humans. In this work we will present methods belonging to two sub areas of computer science: machine learning (ML) and computer vision (including image processing with deep learning). This work is conceived during a time when artificial intelligence in being integrated into processes across different industries and services with the purpose of aiding humans in their activity. Computer vision (classification and detection of objects from images/video) and machine learning are evolving at a rapid pace and provide more and more assistance across many domains. Employing these techniques to optimize, accelerate, simplify or aid either production processes from agriculture or industry or processes related to healthcare is a real necessity, as the demand for them is constantly growing.

Our focus in this work falls on researching how to improve computer vision and machine learning techniques when applied with the purpose of automating existing processes. The two principal fields that we have explored are agriculture and medicine, two unrelated fields, showcasing the flexibility and adaptability of computer vision and machine learning algorithms.

Machine Learning and Computer Vision for Smart Agriculture - we chose this subject as agriculture is a crucial domain for global food supply. As the global population continues to rise and the amount of fertile soil is limited, any innovation that helps to increase the amount of crops/fruits harvested constitutes a valuable addition. Smart agriculture itself, at the time when the thesis was developed, is a field that is not extensively explored, yet holds several directions of research with potential, such as automatic harvest estimation and leaf disease detection from images [JHB19, CSD+17, MLM+20, YWBZ19, EC19]. The former can be employed for early yield estimation, giving more time to farmers to allocate an adequate amount of resources for the actual harvesting process and for storage. The latter can be employed for early-warning systems that can notify farmers in the event that there are sudden changes in the leaf coloration of plants, potentially indicating some disease. Both of these systems enable the optimization of processes in agriculture which in turn helps increase the yield.

Machine Learning and Computer Vision for Medical Diagnosis - quality medical services are among the cornerstones of modern civilisation, as such, it is natural to attempt to optimize and improve them. As previously mentioned, computer science and artificial intelligence have been intertwined with many fields, healthcare being no exception. There is great interest in the research community for computer-aided diagnosis and medical data processing [RIZ+17, RvKF12]. Diagnosis-wise, one major problem is that the demand for highly trained personnel exceeds the supply. Thus, automating partially or fully the process could alleviate part of the demand. Another aspect that poses great research interest is the substitution of invasive diagnosis processes with non-invasive ones [RGBM17], such as using only medical images. In this area, some models already outperform humans [RIZ+17] when working with radiology images. As such there is great promise in the area of using computer vision techniques for medical diagnosis. Furthermore, the medical system could benefit from an early automated warning system which could identify and schedule people for checkups based on their medical data. Still within the healthcare service's concerns is working with large amounts of data generated by collecting medical histories. An improvement could be derived by analysing the data using unsupervised algorithms to extract strongly correlated features and retaining only one of them in order to reduce the data volume.

Machine Learning and Computer Vision in Collaborative Trans-disciplinary Education - this research area combines artificial intelligence techniques, real-world problems and teaching methodologies. As intelligent algorithms are integrated in more and more domains, computer scientists' ability to understand the needs, requirements and perspectives of researchers from other areas constitutes a significant advantage in model development. As such we deemed that is would be useful to introduce a course focused on attempting to solve real-world issues using artificial intelligence methods while modelling the issues from both the perspective of the computer scientist and from the perspective of an expert from the domain of the problem. This would teach students how to broaden their horizon, improve their creativity and critical thinking when solving problems. A secondary motivation for introducing this course was the improvement of our own perspective related to artificial intelligence techniques techniques, as instead of applying these techniques we would need to explain and present them to others [BM19, BKSD10, LZHL18].

1.2 Objectives

The Ph.D. thesis consists of employing and refining artificial intelligence methods applied to real world problems of classification, detection, pattern recognition and data extraction. All of the original research conducted was done under the supervision of Prof. Dr. Horia F. Pop. The aim of this thesis is proposing solutions that automate, either fully or partially, complex real-world tasks. This would improve cost-efficiency and productivity of vital processes that require identifying key elements or entities in structured or unstructured data (images, text). Based on our literature review, ML algorithms and deep learning for image processing techniques are good choices for our objectives since on a high level, these methods are able to extract features required to make decisions that can assist humans and artificial intelligence based solutions can assist the processes while being less expensive to scale up.

The two main fields that were selected for our research were **Agriculture** and **Healthcare** for their importance in the real world and the practical problems that they contain. We have also conducted research that, while not central to the scope of the thesis, benefits from the results and knowledge gained from those experiments, in the form of teaching the application of artificial intelligence (ML, deep learning, etc.) algorithms and data-centric approaches in solving real world problems to college students.

Smart agriculture, based on our literature review, poses several issues and we have formulated objectives centered around them. The most prevalent problem was the lack of publicly available, high quality, annotated datasets with fruit images. Another issue was the difficulty of comparing results obtained in Smart Agriculture, as the datasets used were distinct and the works focused on different tasks. Regarding text data, such as historical information pertaining to seeding dates, locations or crop type, while there are more such datasets, they may contain missing values. As such, we aimed to address these issues. The first objective was to gather, centralize and summarize the state of the art in Smart Agriculture. Following, the next objective was to introduce a new, high quality dataset with images of fruits, from a variety of species, suited for classification tasks. Besides introducing the dataset, we aimed to introduce a model trained and tested on the dataset to establish a baseline performance. Further, we planned to expand the dataset with images and annotation files suited for detection via an algorithm that generates the images based on the classification dataset. An objective linked with the detection dataset was the introduction of a detector trained on purely synthetic images and compare its performance with state of the art detectors trained on real world images in order to reduce the effort dedicated to data acquisition. For text data, we studied imputation methods for handling datasets with missing values in order to improve a model that estimates the yield based on planting date. We also evaluated the usage of unsupervised outlier removal models to further boost an estimator's performance.

For **Healthcare** we analysed the state of the art and selected 2 main objectives. Firstly, with the advent of dermoscopy, a non-invasive skin lesion diagnosis method, we aimed to train a classifier that can be used to identify multiple types of skin lesions. We further expanded on this research by studying what was the impact of oversampling techniques on the performance of the classifier and if these techniques were sufficient to address the issue of underrepresented classes. A second objective was the development of a machine learning model that estimates the risk of osteoporosis in patients based on their medical history. A sub-goal of this objective was improving the retrieval times of data for any given patient, identifying non-trivial data correlations and integrating these functionalities in a software tool usable by medical specialists.

Another objective set was the introduction of a course focused on Machine Learning and Advanced Computer Vision Techniques targeted towards undergraduates. We aimed to form young researchers that are able to approach such problems with a broader perspective, taking into account the interdisciplinary nature of the tasks. Based on our observations from the literature, we also noted that data processing and understanding is a key element to developing viable models for real-world problems, idea that we aimed to pass on to the students. We proposed to cover topics starting from problem modelling, dataset acquisition, exploration, sanitizing, augmentation and processing, model selection, training and testing and result analysis and understanding.

Chapter 2

Machine Learning and Computer Vision for Smart Agriculture

2.1 Background for Smart Agriculture

In this Section we present the background and state of the art contributions that we have leveraged in our research related to fruit classification and detection as well as plant disease identification using deep learning. We have reviewed these papers in order to aid us in our original contributions detailed in [MO18, MCC20a, MCC20b, Mur22]. There is interest in developing classifiers and detectors using images of fruits/plants and regressors using text data with the purpose of optimizing agriculture-related processes, such as harvesting, resource allotment, treatment distribution. We highlight the related problems, such as: *fruit recognition from images using deep learning*, *detecting leaf plant diseases from images*, *deep learning methods for fruit yielding estimation*. Within this Section we will present each problem's relevance in the field and the reasoning for the deep learning approaches.

2.2 Proposed Approaches for Smart Agriculture

This section details our approaches concerning fruit and leaf disease classification, fruit detection from images and yield estimation based on seeding date and weather parameters. We note that the most prevalent issue in this field is the lack of available data. Thus, in our approaches we aimed to address this issue by introducing a large dataset with high quality images containing 90483 images of 131 species of fruit. We describe the process employed to create the fruit images, we also study the viability of using the dataset in conjunction with a deep learning model (based on VGG16 [SZ14]). We compare how small changes in the model's architecture and using different colorspaces can affect the performance and we obtained an accuracy of 98.66%. Compared to the literature, a similar classifier is presented in [XWL+19], obtaining and accuracy of 89.06%. Further, we present how we modelled the issue of creating images viable for detection. We aimed to generate synthetic images, based on the ones we already collected for classification and combine them with background images. This approach was not attempted before in tasks concerning outdoor locations. We describe the algorithm that we designed and implemented, the structure of the detector that we used to test the quality of the generated images and we compare our results with other detectors that were trained solely on real world data. Our best result, in terms of F1-Score (0.816) ranks within the same performance interval as other detectors (other works that we analysed fall within the (0.783, 0.849) interval of F1-Scores). Using information regarding the seeding date and weather parameters, we trained a regression model to identify any patterns for a better yield estimation. We studied the effectiveness of nine different methods for missing value imputation: mean values, similar values, linear, spline and pad interpolation, linear regression, random forest, extreme and histogram gradient boosting regression. Further, we studied outlier removal using unsupervised methods such as isolation forest, minimum covariance determinant, local outlier factor and one class SVM. We achieved an increase in performance for our estimator (measured using the R^2 metric) from 0.723 (in the case of removing samples with missing attributes) to 0.938 when using random forest and one class SVM. We present our conclusions related to both the dataset creation processes and the deep learning models that we have trained and we follow up with future plans for research.

2.3 Chapter conclusions

Based on our research in smart agriculture, on the topic of harvest estimation, we have introduced a dataset with fruit images, Fruits-360, well suited for classification. We have further expanded the usability of the Fruits-360 dataset by introducing a software that generates synthetic images and annotation files which can be used for detection tasks. We introduced a convolutional neural network that was trained on the Fruits-360 dataset and tested it, achieving an accuracy in classification of 98.66%. For the data generation algorithm, we studied if the images being created can be successfully utilised to train a detector and tested it on real world data. Based on the F1-Score of 0.816 and compared to other works, have shown that it is feasible to use generated data as a substitute for real data, without a major penalty to performance. Using historical yield data and meteorological conditions at the time of seeding can be successfully used to create accurate estimators. We further shown that imputation methods and outlier removal should be performed in order to gain a significant boost in performance of the estimators.

In conclusion, our results make researching Smart Agriculture more accessible by alleviating the need to collect real world images and manually annotating them. Furthermore, performing comparative analysis on different detectors can be done more easily, as data can be generated for any fruit detection scenario. As we have shown that training on synthetic data produces satisfactory results in fruit detection, other domains can be approached in a similar manner, inspired by these results.

2.4 Contribution summary

- The introduction of a high quality dataset (Fruits-360) of fruit images for classification (90483 images distributed in 131 species of fruit) constitutes an original contribution as no such public dataset has been released [MO18].

- A novel model for fruit classification that is capable of distinguishing from 131 types of fruits [MO18].

- Addressing the issue of data scarcity for fruit detection by introducing an algorithm that builds a detection dataset using images from the Fruits-360 dataset and creates annotation files. The generated dataset can be customised by the user to fit to their own particular fruit detection scenario. The work constitutes an original contribution as there are no approaches that generate synthetic datasets taking into account outdoor settings [Mur22].

- We introduced a convolutional neural detector network trained on synthetically generated data and obtained comparable results with other state of the art models trained exclusively on real world data. No other works make use of models trained on just synthetic data and tested on real world images [Mur22].

- Trained a regression model for sunflower yield estimation based on seeding date and weather parameters. We studied the effectiveness of multiple methods for missing value imputation and we studied outlier removal using unsupervised methods. We achieved significant increases in the performance of the estimator compared to training on only the data that is not missing any attributes [CCM23].

Chapter 3

Machine Learning and Computer Vision for Medical Diagnosis

3.1 Background for Medical Diagnosis

In this Section we present the background for skin lesion diagnosis and osteoporosis and fractures caused by bone-frailty risk estimation that we have utilised in our research. This knowledge was gathered in order to assist us in the original work done in [Mur19,MMC+21]. This section details the current state of the art in computer aided medical diagnosis. Further, we present the importance of dermoscopy, a non-invasive skin lesion diagnosing technique based on medical images. The current state of the art works regarding skin lesion classification using deep learning models obtains accuracies of 81.33% [RGBM17] and 78% [MdS18] however the main focus falls only on malignant lesions. We highlight as well the need for fracture risk estimation in order to enroll high risk patients in prevention programs. Tools that calculate a risk fracture based on attributes from a person's medical history have been created, such as FRAX and XRAIT [KBB+21] with 70% sensitivity and 92% specificity. Despite this, no tool that calculates this risk factor in real time has been found.

3.2 Proposed Approaches for Medical Diagnosis

In this section we start by presenting our approach to the task of skin lesion classification. The focus in the current state of the art was on distinguishing malignant skin lesions from one another however, in the real world, the predominant class of lesions is benign. Thus, in order to create a model that can provide aid in the general case, we opted to use the HAM-10000 dataset [TRK18], as it includes benign lesions. This adds another layer of complexity, as benign lesions can be visually very similar to malignant lesions. We then discuss the model that we selected for the task, and the methodology employed to address the unbalanced data in the set. Further, we discuss our results (83.96% accuracy) and compare them to those obtained in the literature. Following we present the work done in collaboration with medical specialists from the National Health Service of Scotland focused on analysing anonymised medical data and extracting information that could be used to automatically detect patients with a high risk of fractures. We then showcase a tool that we have developed in parallel to our research that aids patient data retrieval. We summarize our findings from the provided data and we formulate proposals of enhancement.

3.3 Chapter conclusions

Following this research, we have introduced a model based on the Inception-ResNet architecture capable of classifying skin lesions from medical images. This impacts computer aided dermoscopy, increasing its viability as a replacement for the current invasive diagnosis process. We also performed a comparative study of the effect of random oversampling and ADASYN on the performance of the proposed model. We have concluded that, while both oversampling techniques improve the performance of the model, the ADASYN algorithm provides a larger benefit, surpassing similar work presented in [RGBM17] (81.33%) and [MdS18] (78%).

Based on our analysis of the data and the needs of the clinical team, we have developed an proof of concept application that employs fast data retrieval and data presentation to aid clinical decisions. The patient data is retrieved in under two seconds. Clustering shows promising results, however the available data contains purely true positives (ie. patients that have suffered fractures) and, in the absence of true negative data, a robust model cannot be created. In a follow-up research phase, such data will be made available and machine learning algorithms can be trained to identify candidates for the Fracture Liaison Service. The clinical process would permit an intelligent algorithm to run periodically and produce a list of patients sorted in accordance to a risk score. The medical records of those patients can then be analysed by a specialist using the developed interface that presents the most relevant medical information about the patient. Thus we aim to reduce the 6 week delay in patient identification that occurs due to manual data processing.

3.4 Contribution summary

- We introduce novel approaches to skin lesion classification using deep learning (model based on Inception-ResNet) and dermoscopy, a non-invasive diagnosis technique. One element of originality stems from including benign lesions (via the AM-10000 [TRK18] dataset) in our experiments, in order to test our model in conditions that are closer to real world situations. We studied the impact of sampling methods (simple oversampling and ADASYN [HBAGL08]) to handle imbalanced datasets [Mur19].

- We present an analysis of the historical medical data that is linked with bone structure afflictions, data provided by the National Health System Scotland that has been anonymised. The features with the highest impact on the fracture risk are isolated to be used in a model that can automatically estimate this risk [MMC+21].

- We introduce a novel tool, constituting an original contribution, to be used be the clinical team for fast medical history retrieval for any given patient, accelerating the speed at which they can select good candidates to enroll in the Fracture Liaison Service [MMC+21].

Chapter 4

Machine Learning and Computer Vision in Trans-disciplinary Education

4.1 Context for Collaborative Trans-disciplinary Education

In this chapter we present research that, although not belonging to the central investigation of the thesis, has shown great potential and has been an opportunity to transfer the knowledge and insight gained in our other experiments and analysis. Besides our experiments regarding the applications of artificial intelligence techniques for agriculture and healthcare, we aimed to teach and train college students how to better integrate artificial intelligence algorithms in solving complex real-world problems. The majority of state of the art papers are focused on identifying small improvements to the models they employ however, the quality of the training/validation/testing data is highly important and is much less tackled. Obtaining or creating a high quality dataset is a key factor in introducing a high performing model, task which can be time consuming and difficult. Furthermore, machine learning and computer vision are applied in conjunction with one or more fields, such as agriculture, medicine, social sciences, physics, etc. Understanding how to measure the quality of a dataset is an aspect that requires an understanding of the field in which the model is to be used. Thus, we introduce a course that encourages students approach multi-disciplinary problems taking into account information from multiple perspectives, to re-frame these problems such that they can be tackled using artificial intelligence methods and to understand how to analyze, model and process data coming from other fields. We will detail the course structure, the main ideas that we aim to teach, the evaluation method and the feedback that we have received after this course. This not only serves as means to prepare future generations of researchers but also as means to strengthen our own grasp on trans-disciplinary and inter-disciplinary approaches to problem solving.

4.2 Course details

The course combines several techniques from the field of artificial intelligence with expert knowledge from the climate change domain, solidifying the multi and trans-disciplinary aspects of the course. Another novel characteristic of the course is the modern approach towards collaborative and experiential learning for students, teachers and external specialists. The evaluation methodology, also reflecting of the collaborative approach, involves self-evaluation and inter-evaluation. The quality of the solutions for problems pertaining to climate change has been linked to the application of multi-disciplinary techniques in creating these solutions however, the management of these methods is a key factor in the probability of obtaining a good solution [BM19], thus collaboration and experimentation play a large role.

The projects that have been worked on tackled topics such as: predicting drought and floods, classifying and predicting the evolution of air quality, fruit classification and detection from images, animal detection from images, detecting storm formations, monitoring deforestation, predicting the occurrence of forest fires, smart plant irrigation, automatic monitoring of greenhouses, estimation of weather and temperature evolution, analysis of city metabolism, classification of care homes and orphanages based on needs and priorities, monitoring household electricity consumption, predicting the traffic in big cities applied to the big cities from Europe. The applied methods include Long short-term memory, K-Means clustering, Principal component analysis, neural networks (regressors, classifiers; with and without convolutions) such as DeepWeeds, Inception-v3, ResNet-50, Support vector machines, Random forest. All teams were able to source appropriate training and testing data using websites such as Kaggle, Google Datasets or by contacting authors of datasets that are not freely accessible. Further, all teams were able to setup pipelines for training, validating and testing their models of choice, obtaining explainable results.

4.3 Chapter conclusions

The course approach that was introduced was new in the context of the faculty and was a new experience for the students. Their feedback and the interest shown by the external guests indicate that the methods could achieve even greater success in the future iterations. To the extent of our knowledge, this method of course teaching is unique in the computer science faculties of Romania. At an international level, this method has been seen in courses like "AI for Social Good" from Stanford University, "Climate Change and AI" from the University of Oxford, and "AI for the Study of Environmental Risks" from the University of Cambridge.

4.4 Contribution summary

- We introduced a new course focused on a variety of artificial intelligence techniques (such as ML and deep learning) applied to subjects that require knowledge from other disciplines such as physics, biology or chemistry. The original contribution consists of collaborating with multiple external specialists to organise lectures with the purpose of presenting a different perspective to the problems being tackled by the enrolled students [CCM22].

Chapter 5

Concluzii

We have studied and applied several deep learning models based on architectures such as VGG [SZ14], Residual Network [HZRS15] and SSD [LAE+15] and we have applied them to the task of fruit classification and detection based on images taken in orchard and plantation environments. We have addressed the major issue of data scarcity in the area of Smart Agriculture by introducing the Fruits-360 dataset [OMa, OMb] containing 131 types of fruit totalling 90483 images for the purpose of classification. We have expanded the usability of the dataset by the addition of an algorithm that generates synthetic fruit images and bounding box annotations, which is well suited for detection tasks. We have introduced a model based on the SSD architecture for detecting apples from images trained exclusively on synthetic images and tested on real-world images. The performance obtained was on par with similar works that have utilised solely real-world images. Our proposed methods have been validated by publications in reputable conferences, demonstrating that our approach is useful and of interest for both the artificial intelligence and the smart agriculture communities.

Regarding our results, we have reviewed and aggregated multiple state of the art works in the area of fruit detection and leaf disease identification for a better understanding of the existing challenges in this domain. Our reviews have been published in [MCC20a, MCC20b]. We have introduced the Fruit-360 dataset and we have established a baseline performance for classifiers using a model based on the VGG16 architecture, which was published in [MO18]. Lastly, we have introduced a method of generating synthetic fruit images with annotated bounding boxes for detection, we have tested the performance that can be obtained by utilising only synthetic data for training and achieved results in the same performance range as projects using exclusively real-world data. We have published these results in [Mur22].

Our future work will focus on further refining the data generation algorithm in order to create more natural looking images which could further improve the performance of the trained detectors. One possibility is the usage of Generative Adversarial Networks trained on fruit images from the Fruits-360 dataset and applied on images with orchard-like backgrounds.

For the field of medical diagnosis, we have presented a model based on Inception-ResNet [SIV16] capable of classifying seven types of skin lesions from medical images. We have analysed the impact of applying oversampling algorithms on image datasets on the performance of the trained classifiers. We have also approached the issue of estimating the risk of osteoporosis and frailty fracture occurrence using multiple machine learning algorithms on data provided by NHS Scotland. Our proposed methods have been validated by publications in reputable conferences, demonstrating that our approach is useful and of interest for the machine learning and the medical communities.

The results related to skin lesion classification were published in [Mur19]. We have shown that employing an adaptive synthetic sampling method on skin lesion images the performance of the classifier improves, while also obtaining better results than previous state of the art works. For osteoporosis risk evaluation we have applied machine learning algorithms to select a subset of all the features collected by the medical specialists and identify patterns in a large volume of data. We have also implemented a fast data retrieval algorithm that aids medical staff in consulting historical patient data. This research has been published in [MMC+21].

In the future we plan on expanding the research by integrating texture analysis in the lesion classification task and study the impact of this addition when attempting to separate between melanoma and melanocytic nevi. For osteoporosis risk estimation, we will focus on selecting a subset of features that are not strongly correlated to one another and design an artificial intelligence model that can calculate a risk score to aid clinicians. For this however, we are dependant on more medical data to be provided in order to create a solid training set.

Bibliography

[BKSD10] Marlies Baeten, Eva Kyndt, Katrien Struyven, and Filip Dochy. Using student-centred learning environments to stimulate deep approaches to learning: Factors encouraging or discouraging their effectiveness. *Educational Research Review*, 5(3):243–260, 2010.

[BM19] Lino Briguglio and Stefano Moncada. The benefits and downsides of multidisciplinary education relating to climate change. *Climate Change and the Role of Education*, pages 169–187, 2019.

[CCM22] Adriana Mihaela Coroiu, Alina Delia Călin, and Horea-Bogdan Mureşan. Collaborative transdisciplinary educational approaches in ai. In *International Conference on Computer Supported Education CSEDU (2)*, pages 260–267, 2022.

[CCM23] Alina Delia Călin, Adriana Mihaela Coroiu, and Horea Bogdan Mureşan. Analysis of preprocessing techniques for missing data in the prediction of sunflower yield in response to the effects of climate change. *Applied Sciences*, 13(13), 2023.

[CSD+17] Steven W Chen, Shreyas S Shivakumar, Sandeep Dcunha, Jnaneshwar Das, Edidiong Okon, Chao Qu, Camillo J Taylor, and Vijay Kumar. Counting apples and oranges with deep learning: A data-driven approach. *IEEE Robotics and Automation Letters*, 2(2):781–788, 2017.

[EC19] Ngozi Clara Eli-Chukwu. Applications of artificial intelligence in agriculture: A review. *Engineering, Technology & Applied Science Research*, 9(4):4377–4383, 2019.

[HBAGL08] Haibo He, Yang Bai, Edwardo A. Garcia, and Shutao Li. Adasyn: Adaptive synthetic sampling approach for imbalanced learning. In *IJCNN*, pages 1322 – 1328, 07 2008.

[HZRS15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015.

[JHB19] F Jakjoud, A Hatim, and A Bouaaddi. Deep learning application for plant diseases detection. In *Proceedings of the 4th International Conference on Big Data and Internet of Things*, pages 1–6, 2019.

[KBB+21] Nithin Kolanu, A Shane Brown, Amanda Beech, Jacqueline R Center, and Christopher P White. Natural language processing of radiology reports for the identification of patients with fracture. *Archives of Osteoporosis*, 16(1):1–8, 2021.

[LAE+15] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott E. Reed, Cheng-Yang Fu, and Alexander C. Berg. SSD: single shot multibox detector. *CoRR*, abs/1512.02325, 2015.

[LZHL18] Bin Liu, Yun Zhang, DongJian He, and Yuxiang Li. Identification of apple leaf diseases based on deep convolutional neural networks. *Symmetry*, 10(1):11, 2018.

[MCC20a] Horea-Bogdan Mureşan, Alina Delia Călin, and Adriana Mihaela Coroiu. Overview of recent deep learning methods applied in fruit counting for yield estimation. *Studia Universitatis Babeș-Bolyai Informatica*, 65(2):50–65, 2020.

[MCC20b] Horea-Bogdan Mureşan, Adriana Mihaela Coroiu, and Alina Delia Călin. Detecting leaf plant diseases using deep learning: A review. In *2020 International Conference on Software, Telecommunications and Computer Networks (SoftCOM)*, pages 1–6, 2020.

[MdS18] Danilo Barros Mendes and Nilton Correia da Silva. Skin lesions classification using convolutional neural networks in clinical images. *CoRR*, abs/1812.02316, 2018.

[MLM+20] Shihan Mao, Yuhua Li, You Ma, Baohua Zhang, Jun Zhou, and Kai Wang. Automatic cucumber recognition algorithm for harvesting robots in the natural environment using deep learning and multi-feature fusion. *Computers and Electronics in Agriculture*, 170:105254, 2020.

[MMC+21] Alexandru-Ion Marinescu, Horea-Bogdan Mureşan, Alina-Delia Călin, Adriana-Mihaela Coroiu, and Maria Talla. Freida - fracture risk evaluation using highly efficient information retrieval and analysis of large healthcare datasets. In *2021 IEEE International Conference on Big Data (Big Data)*, pages 3917–3922, 2021.

[MO18] Horea Mureşan and Mihai Oltean. Fruit recognition from images using deep learning. *Acta Universitatis Sapientiae, Informatica*, 10(1):26 – 42, 2018.

[Mur19] Horea-Bogdan Mureşan. Skin lesion diagnosis using deep learning. In *2019 IEEE 15th International Conference on Intelligent Computer Communication and Processing (ICCP)*, pages 499–506, 2019.

[Mur22] Horea-Bogdan Mureşan. An automated algorithm for fruit image dataset building. In *2022 17th Conference on Computer Science and Intelligence Systems (FedCSIS)*, pages 103–107, 2022.

[OMa] Mihai Oltean and Horea Muresan. Fruits 360 dataset on github. `https://github.com/Horea94/Fruit-Images-Dataset`. [Online; accessed 15.06.2021].

[OMb] Mihai Oltean and Horea Muresan. Fruits 360 dataset on kaggle. `https://www.kaggle.com/moltean/fruits`. [Online; accessed 15.06.2021].

[RGBM17] A. Romero Lopez, X. Giro-i-Nieto, J. Burdick, and O. Marques. Skin lesion classification from dermoscopic images using deep learning techniques. In *2017 13th IASTED International Conference on Biomedical Engineering (BioMed)*, pages 49–54, Feb 2017.

[RIZ+17] Pranav Rajpurkar, Jeremy Irvin, Kaylie Zhu, Brandon Yang, Hershel Mehta, Tony Duan, Daisy Ding, Aarti Bagul, Curtis Langlotz, Katie Shpanskaya, Matthew P. Lungren, and Andrew Y. Ng. Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. *CoRR*, abs/1711.05225, 2017.

[RvKF12] Dirk J Ruiter, Marlieke TR van Kesteren, and Guillen Fernandez. How to achieve synergy between medical education and cognitive neuroscience?

an exercise on prior knowledge in understanding. *Advances in health sciences education*, 17(2):225–240, 2012.

[SIV16] Christian Szegedy, Sergey Ioffe, and Vincent Vanhoucke. Inception-v4, inception-resnet and the impact of residual connections on learning. *CoRR*, abs/1602.07261, 2016.

[SZ14] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014.

[TRK18] Philipp Tschandl, Cliff Rosendahl, and Harald Kittler. The HAM10000 dataset: A large collection of multi-source dermatoscopic images of common pigmented skin lesions. *CoRR*, abs/1803.10417, 2018.

[XWL+19] Qian Xiang, Xiaodan Wang, Rui Li, Guoling Zhang, Jie Lai, and Qingshuang Hu. Fruit image classification based on mobilenetv2 with transfer learning technique. In *Proceedings of the 3rd International Conference on Computer Science and Application Engineering*, pages 1–7, 2019.

[YWBZ19] Rongjian Yang, Minghui Wu, Zhaohua Bao, and Pei Zhang. Cherry recognition based on color channel transform. In *Proceedings of the 2019 International Conference on Artificial Intelligence and Computer Science*, pages 292–296, 2019.