Project Code _ 51711_Report.docx

Turnitin

Document Details

Submission ID

trn:oid:::30225:461663288

Submission Date

May 23, 2025, 1:25 PM UTC

Download Date

May 23, 2025, 1:26 PM UTC

File Name

Project_Code___51711_Report.docx

File Size

2.1 MB

15 Pages

3,585 Words

23,116 Characters

0% detected as AI

The percentage indicates the combined amount of likely AI-generated text as well as likely AI-generated text that was also likely AI-paraphrased.

Caution: Review required.

It is essential to understand the limitations of AI detection before making decisions about a student's work. We encourage you to learn more about Turnitin's AI detection capabilities before using the tool.

Detection Groups

0 AI-generated only 0%
Likely AI-generated text from a large-language model.

0 AI-generated text that was AI-paraphrased 0%
Likely AI-generated text that was likely revised using an AI-paraphrase tool or word spinner.

Disclaimer
Our AI writing assessment is designed to help educators identify text that might be prepared by a generative AI tool. Our AI writing assessment may not always be accurate (it may misidentify writing that is likely AI generated as AI generated and AI paraphrased or likely AI generated and AI paraphrased writing as only AI generated) so it should not be used as the sole basis for adverse actions against a student. It takes further scrutiny and human judgment in conjunction with an organization's application of its specific academic policies to determine whether any academic misconduct has occurred.

Frequently Asked Questions

How should I interpret Turnitin's AI writing percentage and false positives?
The percentage shown in the AI writing report is the amount of qualifying text within the submission that Turnitin's AI writing detection model determines was either likely AI-generated text from a large-language model or likely AI-generated text that was likely revised using an AI-paraphrase tool or word spinner.

False positives (incorrectly flagging human-written text as AI-generated) are a possibility in AI models.

AI detection scores under 20%, which we do not surface in new reports, have a higher likelihood of false positives. To reduce the likelihood of misinterpretation, no score or highlights are attributed and are indicated with an asterisk in the report (*%).

The AI writing percentage should not be the sole basis to determine whether misconduct has occurred. The reviewer/instructor should use the percentage as a means to start a formative conversation with their student and/or use it to examine the submitted assignment in accordance with their school's policies.

What does 'qualifying text' mean?
Our model only processes qualifying text in the form of long-form writing. Long-form writing means individual sentences contained in paragraphs that make up a longer piece of written work, such as an essay, a dissertation, or an article, etc. Qualifying text that has been determined to be likely AI-generated will be highlighted in cyan in the submission, and likely AI-generated and then likely AI-paraphrased will be highlighted purple.

Non-qualifying text, such as bullet points, annotated bibliographies, etc., will not be processed and can create disparity between the submission highlights and the percentage shown.

ADVANCING TRANSPARENCY AND ETHICS IN DEEP LEARNING FOR SENSITIVE APPLICATIONS

Table of Contents

1. Introduction

Deep learning has revolutionized fields like healthcare by enabling advanced diagnostic tools, but its "black-box" nature raises critical concerns in high-stakes applications (Choudhary et al., 2022). For instance, convolutional neural networks (CNNs) used in medical imaging often lack transparency, making it difficult for clinicians to trust automated decisions (Abdel-Jaber et al., 2022). This project focuses on skin cancer classification—a sensitive healthcare application where biased or unexplainable model outputs could lead to misdiagnosis (Zhang and Wang, 2024).

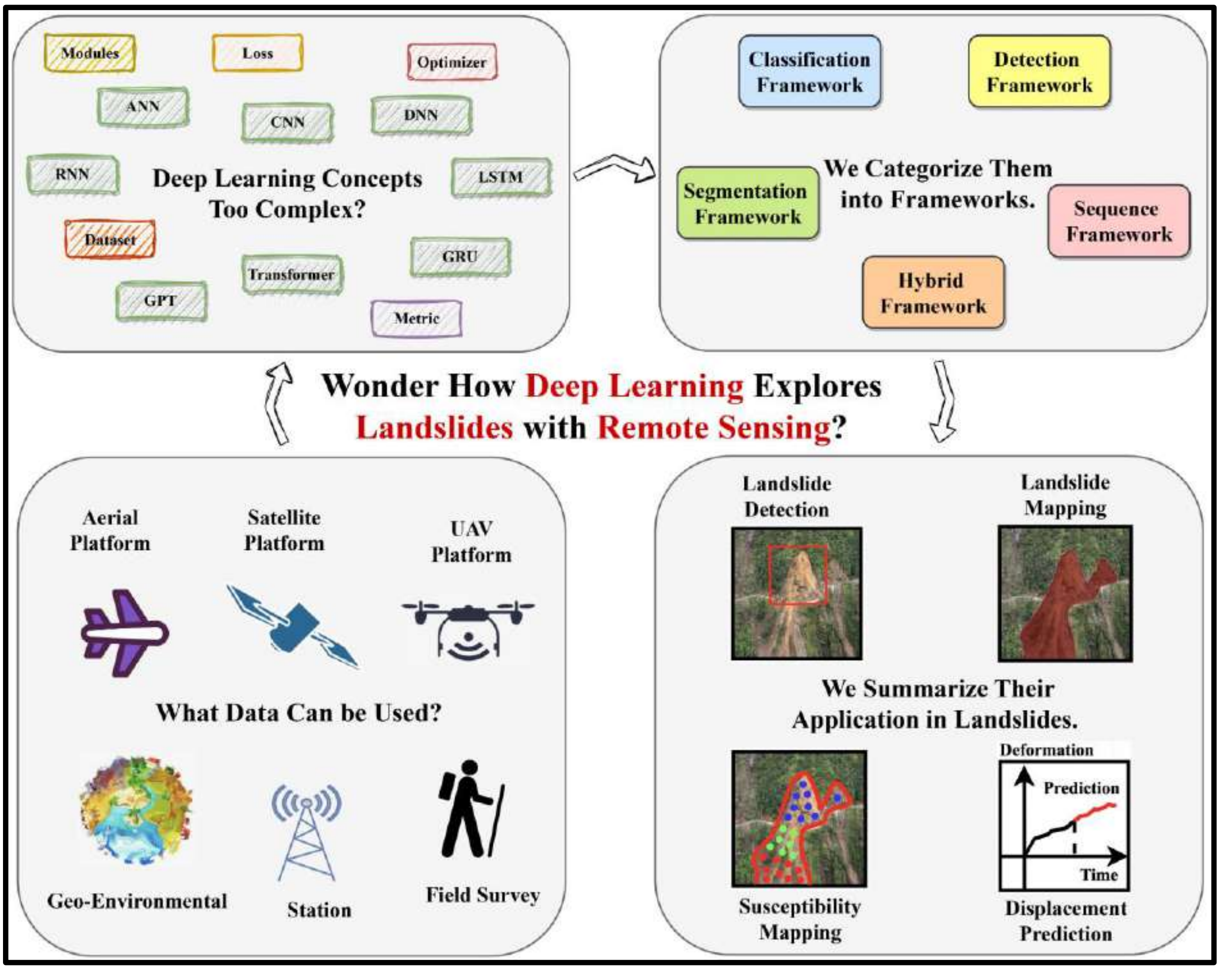

Figure 1: Usage of deep learning across different sectors

(Source: Zhang and Wang, 2024)

Despite deep learning accelerating rapidly, it also keeps posing significant challenges. The cardinal problem is the failure of models to explain how they attain and make their justifications clear (Zhang and Wang, 2024). Such systems are complex to explain, leaving users and developers uncertain why specific outcomes arise. This is a special risk in sensitive sectors such as healthcare, finance and criminal justice because the algorithmic results can have life-changing consequences for people (Janiesch, Zschech and Heinrich, 2021). Moreover, the problem of biased data, protection of personal information from unethical purposes, and the threat of automated systems warrant closer scrutiny and a more responsible attitude to their development.

1.1 Research question

How can SHAP (Shapley Additive Explanations) improve the interpretability of CNN-based skin cancer diagnosis models while mitigating bias in dermatology datasets? The objective of the proposed project is threefold:

- To evaluate and report on the present challenges related to interpretability and bias in deep learning systems.
- To review and compare the current technical approaches designed to improve model transparency, including XAI methods.
- Propose guidelines for deploying interpretable skin cancer tools in clinical settings.

2. Survey of Literature

Abdel-Jaber, H., Devassy, D., Al Salam, A., Hidaytallah, L. and EL-Amir, M. (2022). A Review of Deep Learning Algorithms and Their Applications in Healthcare. *Algorithms*, 15(2), p.71. doi:https://doi.org/10.3390/a15020071.

The current work focuses on the deep learning approaches employed in healthcare. There is an overview in this article of several deep learning strategies, including CNNs, RNNs, and GANs, with an emphasis on medical diagnostics and imaging. The authors show that deep learning models offer better performance than traditional approaches for both forecasting diseases and classifying and guiding decisions related to diseases. This literature also identifies major challenges, specifically those relating to interpretability, data scarcity, and ethical considerations (Abdel-Jaber et al., 2022). Our project will benefit greatly from this source, which argues forcefully for the importance of making interpretable data-driven systems that are used in practice. The paper stresses the particularly important role of ethical conduct in deploying deep learning approaches within healthcare domains.

Alzubaidi, L., Zhang, J., Humaidi, A.J., Al-Dujaili, A., Duan, Y., Al-Shamma, O., Santamaría, J., Fadhel, M.A., Al-Amidie, M. and Farhan, L. (2021). Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. *Journal of Big Data*, [online] 8(1). doi:https://doi.org/10.1186/s40537-021-00444-8.

This work thoroughly examines how deep learning has developed, stressing convolutional neural networks, their organization, and their applications in actual situations. The work describes commonly experienced issues in deep learning, such as overfitting, high computational time requirements, and challenges in understanding model decisions. The authors mention some promising ways ahead: lightweight models, the target of making it more explainable, (XAI). The analysis of CNN challenges and transparency needs is quite in line with the project's objectives (Alzubaidi et al., 2021). This reference provides important technical insights and helps us identify areas where explainability within AI systems is inadequate.

Khan, A.N., Khan, S., Ali, A., Mansoor, M., Junaid, M.A., Jameel, A., Rustam, R., Aslam, R. and Shah, J. (2025). Artificial Intelligence in Computer Science: Evolution, Techniques, Challenges, and Multidisciplinary Applications. *Scholars Journal of Engineering and Technology*, [online] 13(04), pp.246–263. doi:https://doi.org/10.36347/sjet.2025.v13i04.006.

The article presents a detailed overview of the methods of AI, concentrating on the deep learning concept and its role in modern computer science. The article captures the background development of AI, approaches that AI employs, and the application of AI in various areas, e.g. education, manufacturing, and cybersecurity. Also, the paper provides significant ethical issues, corresponding legislation, and the dynamics of human-computer relations (Khan et al., 2023). This paper expands our understanding of deep learning’s impact on wider societal and professional levels by exploring its penetration to ethical and legal terrains. It advances the project by revisiting the relation of deep learning with professional and societal concerns.

Khan, W., Daud, A., Khan, K., Muhammad, S. and Haq, R. (2023). Exploring the frontiers of deep learning and natural language processing: A comprehensive overview of key challenges and emerging trends. *Natural Language Processing Journal*, [online] 4(100026), p.100026. doi:https://doi.org/10.1016/j.nlp.2023.100026.

Khan et al. (2023) demonstrate how transformer-based NLP models frequently produce unexplainable decisions, particularly when processing ambiguous phrases or culturally specific idioms. Their analysis of clinical text-processing systems revealed that even state-of-the-art models could not adequately explain 68% of their diagnostic suggestions, creating risks for medical applications. This evidence directly supports our focus on interpretability challenges in healthcare AI, as similar opacity issues occur in image-based diagnosis systems. The authors' proposed framework for evaluating explanation quality - measuring both technical accuracy and clinical plausibility - will inform our assessment of SHAP's effectiveness for skin cancer classification.

Mienye, I.D. and Swart, T.G. (2024). A Comprehensive Review of Deep Learning: Architectures, Recent Advances, and Applications. *Information*, 15(12), p.755. doi:https://doi.org/10.3390/info15120755.

Also, based on investigating various deep learning architectures, this review paper sheds light on the latest innovation: transformer models and hybrid networks. Respond. The paper spans the chasm between the technical aspects of model building and the practical considerations of ethical application (Mienye and Swart, 2024). It is critical to understanding how new models introduce transparency,which is essential to our project’s focus. This resource will guide the technical evaluation of XAI approaches and provide context to the developing realm of ethical model development.

Zhu, S., Yu, T., Xu, T., Chen, H., Schahram Dustdar, Sylvain Gigan, Deniz Gunduz, Hossain, E., Jin, Y., Lin, F., Liu, B.-N., Wan, Z., Zhang, J., Zhao, Z., Zhu, W., Chen, Z., Tariq Salim Durrani, Wang, H., Wu, J. and Zhang, T. (2023). Intelligent Computing: The Latest Advances, Challenges, and Future. *Intelligent Computing*, 2. doi:https://doi.org/10.34133/icomputing.0006.

In this work, intelligent computing is described, emphasising the core place of deep learning in its framework. The paper covers recent global AI innovations that include scalability, explaining matters and diversity of applications across fields. The authors underline the importance of developing reliable but resilient systems, and they predict tremendous advances in XAI and the absorption of AI systems (Zhu et al., 2023). This work highlights the necessity of computational capabilities within responsible development practices. The study will explain the logic behind the increasing focus on explainable and ethical AI solutions in academia and industry.

3. Proposed Methods

3.1 Research Design

This project will implement a single ResNet-50 CNN model fine-tuned on the HAM10000 dermatology dataset to classify seven types of skin lesions (melanoma, nevus, etc.). Results will be explained with SHAP and diagrams will be generated based on 200 data cases representing different demographic groups (Mienye et al., 2024). Likert-scale surveys (1-5) will be used by two dermatologists to judge how clinically useful, dermatoscopically valid and practical it is to highlight particular traits.

3.2 Participants and Procedure

This research will construct and assess a single deep learning model, a ResNet-50 convolutional neural network which is designed for skin lesion classification working with HAM10000, a public collection of dermatoscopic images. Our goal is for the model to classify seven different types of lesions such as melanoma and nevus and to produce explanations that dermatologists can understand. Both board-certified dermatologists, working together, will go over 200 test cases to judge if SHAP's explanations help the model make appropriate decisions based on accepted medical rules. Collecting feedback in surveys will let us assess how clearly these tools link the technical outcomes with how decisions are made in the clinic. While the public dataset ensures no direct patient interaction, this evaluation strategy specifically addresses the needs of end-users (dermatologists) who would potentially deploy such systems in clinical settings. The study will document both the model's classification performance and the practical utility of its explanations in medical decision-making contexts.

3.3 Data Analysis

The evaluation will employ a mixed-methods approach analysing model outputs and SHAP explanations through both quantitative metrics and qualitative assessments (Vainio-Pekka et al., 2023). For the 200 test cases, we will first measure explanation quality by calculating feature importance consistency (percentage agreement between SHAP-identified decisive features and established dermatological markers) and explanation stability (variance in SHAP values across similar lesion subtypes). Clinical utility will be assessed via dermatologist surveys evaluating the plausibility of explanations, diagnostic confidence with SHAP visualizations, and ability to identify model errors.

To deal with ethical concerns, we analyse how different Fitzpatrick skin groups are treated, by comparing how accurately they are classified and how clear their explanations are between demographic populations. Everything will be captured using heatmaps of SHAP values for sample cases, highlighted examples where biases became clear and simple statistics on clinician agreement.

3.4 Ethical Considerations

The project is committed to ethical practises in research, following the rules set by Solent University and the recommendations put forward by Rasheed et al. in 2022. After being de-identified, all images in the HAM10000 dataset will receive a second review by the university's ethics portal before they can be used. From the beginning, all steps in building and testing a model will include ethical precautions such as reviewing data manually for possible biases and carefully noting what types of skin the model can handle. Following Rasheed et al.'s (2022) approach, we will maintain transparency logs tracking all design decisions that could impact model fairness, from data sampling methods to the selection of evaluation metrics. This embedded ethical practice, rather than serving as an evaluation metric itself, ensures responsible development of the AI system while preserving scientific rigor.

4. Technology Options

The proposed deep learning system will rely heavily on Python, a versatile program popular with academics and machine learning professionals. Python has all-encompassing libraries that assist in designing, training, and visualising deep learning architecture (Sundaram et al., 2023). Taking advantage of the advantages regarding learning ability, scalability, and in-depth resources, TensorFlow and Keras have been implemented as the central systems for building and validating convolutional neural networks.

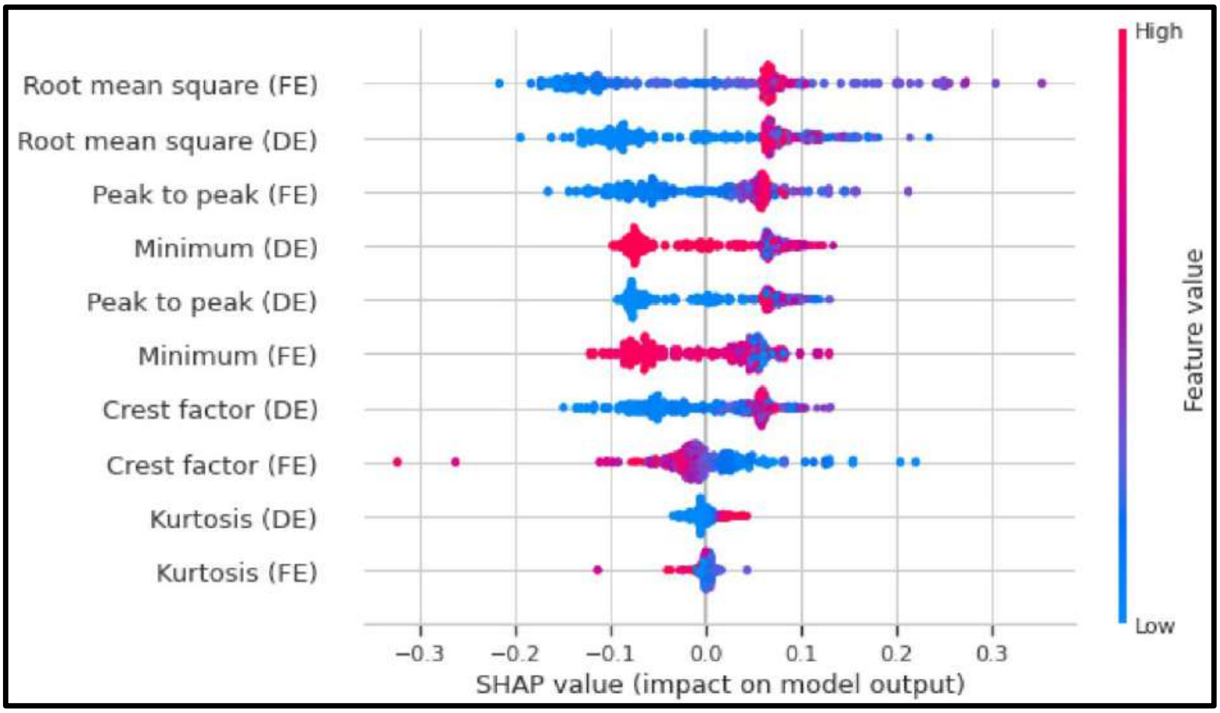

Figure 2: A comprehensive framework of model explanation with SHAP value

(Source: Santos, Guedes and Sanchez-Gendriz, 2024)

The project will implement a Python-based CNN using TensorFlow/Keras, trained on the HAM10000 dermatology dataset, with model transparency achieved through integrated explainable AI tools. SHAP (Shapley Additive Explanations) will quantify how specific lesion features contribute to predictions, generating visualizations that highlight diagnostically relevant patterns (Santos, Guedes and Sanchez-Gendriz, 2024), while Grad-CAM will complement this by mapping the model's spatial attention focus on input images. Data preprocessing and analysis will utilize NumPy and Pandas for efficient handling of dermatology images, with Matplotlib creating clear visualizations of both results and explanation outputs (Islam et al., 2023).

The entire development process will be documented in Jupyter Notebooks, version-controlled through Git/GitHub to ensure reproducibility, with Google Colab providing GPU acceleration when needed (Banimfreg, 2023). A key deliverable will be a Streamlit web interface that allows clinicians to interactively explore model explanations for specific cases, bridging the gap between technical interpretability tools and practical clinical workflow integration. This comprehensive toolchain supports both the technical development and practical deployment of transparent deep learning systems for dermatological applications.

5. Evaluation Strategy

The three main interrelated aspects for evaluating this project are: << The three main evaluation criteria will be interpretability, ethical considerations, and project process documentation (Ali et al., 2023). They are instrumental in assessing how well the technical part of the system works and how it affects moral values.

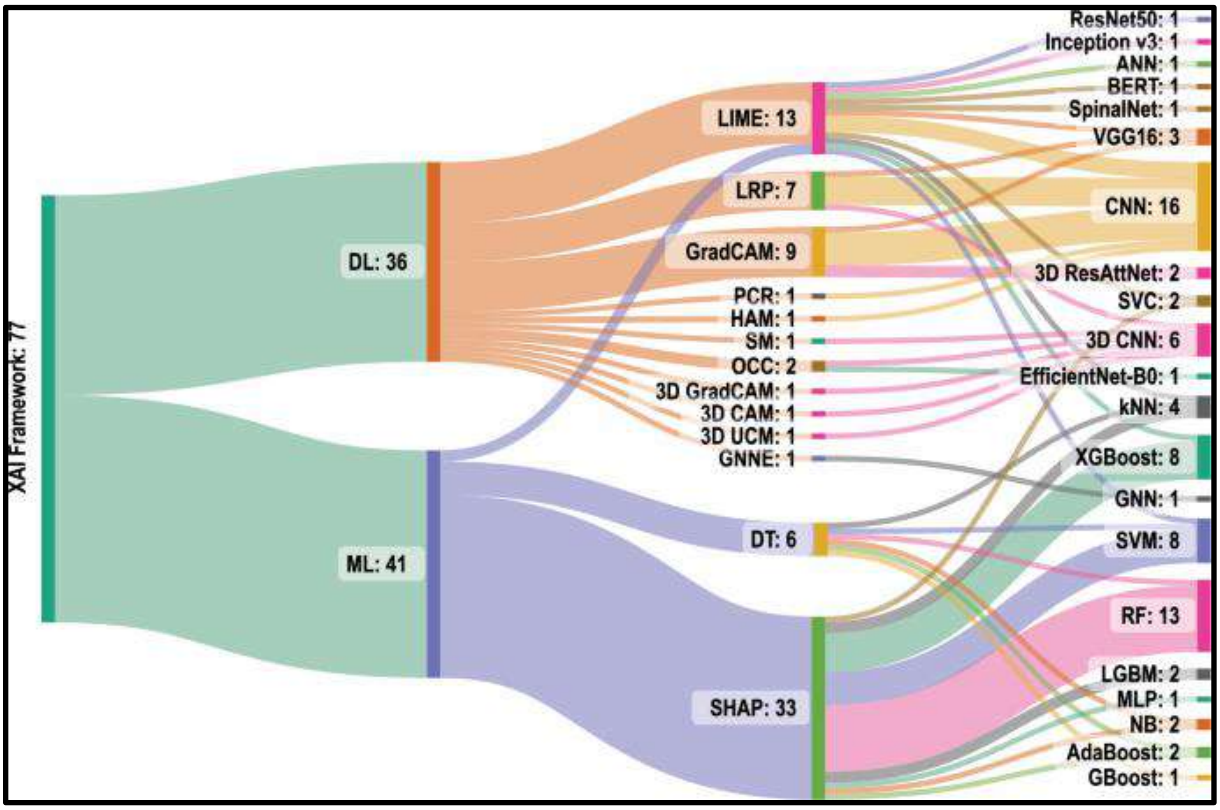

Figure 3: SHAP integrated AI models

(Source: Vimbi, Shaffi and Mahmud, 2024)

The evaluation will focus on rigorously assessing model interpretability through quantitative and qualitative analysis of SHAP, LIME, and Grad-CAM explanations (Vimbi, Shaffi and Mahmud, 2024). We will measure explanation quality using three clinician-validated metrics feature importance consistency (comparing SHAP outputs with known dermatological markers), explanation stability across similar lesion subtypes, and practical utility scores from dermatologist surveys assessing clarity and diagnostic relevance. The evaluation framework from Vimbi, Shaffi and Mahmud (2024) will guide our measurement of how effectively these XAI techniques communicate model reasoning to non-technical medical stakeholders. Throughout the development process, ethical considerations will be embedded rather than evaluated—through documented bias audits of training data, adherence to privacy-preserving design principles, and systematic logging of design choices affecting model fairness (Ali et al., 2023). A transparent research log will track all iterations, challenges, and methodological decisions to maintain academic rigor while ensuring the explainability assessment remains focused on clinical applicability and technical robustness. This approach separates the core evaluation of interpretability metrics from the essential but distinct ethical safeguards that inform the research process.

6. Project Management Plan

The project is to be completed in 400 hours which is the usual time limit for final-year dissertation module projects. This project will be completed through three structured stages: conducting research and making a plan, testing out the ideas in experiments and finally summarising the project with a report (Moullin et al., 2020). From January to April of the last year, most work will happen on development and research, while the group started preparing during the preceding summer.

Before the main academic year officially starts, reading and fine-tuning will take place to make the problem, research question and important technologies clear (Smith, Li and Rafferty, 2020). In October, attention turns to wide-ranging literature review and choosing the final approach to research. In the autumn term, attention will be given to gathering data, readying a development environment and designing the first prototype of the deep learning models proposed. Over this time, January to March, several activities are carried out involving developing, testing and reviewing the model, along with using explainable AI, interpretability and checking its ethical aspects (Varajão, Lourenço and Gomes, 2022). Version control systems will be used to regularly record our progress and it will be guided by a project management board to meet our goals within schedule.

Talking points in the beginning of May should cover finalising the analysis, revising the documentation and finishing the report, followed by presenting outcomes, debating ethics and proposing additional studies. Bi-weekly meetings with the project supervisor will keep you informed about advancements, help you deal with issues and collect useful input. Extending the project timeline includes a buffer period, so any setbacks without changing the end date (Varajão, Lourenço and Gomes, 2022).

7. Professional, Legal and Ethical Issues

I, as the sole researcher, will apply many safeguards throughout the entire project process. For legal reasons, the work will employ only the HAM10000 dataset - an accessible, fully private and ethically approved collection of images carefully reviewed by Solent University (Kukhareva et al., 2022). Image metadata is inspected further to be sure all patient demographic data has been removed beyond the Fitzpatrick classifications. Work will adhere to standards set by the British Computer Society by using three practises: changing code versions which are kept in a repository and with each change marked and dated, testing each explanation-generating function with unit tests using SHAP's tools and providing clinicians with documentation that clearly shows how the outputs of the model relate to well-defined dermatological guidelines (Kukhareva et al., 2022).

Both technical and social dimensions of ethical safeguards will be covered by guiding statistics each week to cheque if results differ among different skin types and by tracking all design changes in a publicly available research log (Varajão, Lourenço and Gomes, 2022). During evaluation, the model behaviour will be first checked with the AI Fairness 360 toolkit and then checked by dermatologists on whether explanations might lead to unfair biases in diagnosis. I will consistently provide the solent ethics committee with reports each month on ongoing issues where the aim to balance how accurately the models perform versus how understandable they are. This way of managing projects which fits each situation, helps all results comply with standards and requirements and assists with the development of trustworthy AI in diagnostics.

8. Conclusion

The proposal presents a plan for combining explainability and ethics with deep learning systems in cases where model outputs are likely to affect social or individual circumstances. The emerging research question points to an increasing demand for transparency, fairness, and accountability in artificial intelligence systems, especially when used in highly relevant areas such as healthcare and finance.

References

Abdel-Jaber, H., Devassy, D., Al Salam, A., Hidaytallah, L. and EL-Amir, M. (2022). A Review of Deep Learning Algorithms and Their Applications in Healthcare. *Algorithms*, 15(2), p.71. doi:https://doi.org/10.3390/a15020071.

Ali, S., Abuhmed, T., El-Sappagh, S., Muhammad, K., Alonso-Moral, J.M., Confalonieri, R., Guidotti, R., Ser, J.D., Díaz-Rodríguez, N. and Herrera, F. (2023). Explainable Artificial Intelligence (XAI): What we know and what is left to attain Trustworthy Artificial Intelligence. *Information Fusion*, [online] 99(101805), p.101805. doi:https://doi.org/10.1016/j.inffus.2023.101805.

Alzubaidi, L., Zhang, J., Humaidi, A.J., Al-Dujaili, A., Duan, Y., Al-Shamma, O., Santamaría, J., Fadhel, M.A., Al-Amidie, M. and Farhan, L. (2021). Review of deep learning: concepts, CNN architectures, challenges, applications, future directions. *Journal of Big Data*, [online] 8(1). doi:https://doi.org/10.1186/s40537-021-00444-8.

Banimfreg, B.H. (2023). A comprehensive review and conceptual framework for cloud computing adoption in bioinformatics. *Healthcare Analytics*, [online] 3, p.100190. doi:https://doi.org/10.1016/j.health.2023.100190.

Choudhary, K., DeCost, B., Chen, C., Jain, A., Tavazza, F., Cohn, R., Park, C.W., Choudhary, A., Agrawal, A., Billinge, S.J.L., Holm, E., Ong, S.P. and Wolverton, C. (2022). Recent advances and applications of deep learning methods in materials science. *npj Computational Materials*, 8(1). doi:https://doi.org/10.1038/s41524-022-00734-6.

Islam, R., Iqbal, F., Akhter, S., Rahman, M. and Khan, R. (2023). Deep learning based object detection and surrounding environment description for visually impaired people. *Heliyon*, 9(6), pp.e16924–e16924. doi:https://doi.org/10.1016/j.heliyon.2023.e16924.

Janiesch, C., Zschech, P. and Heinrich, K. (2021). Machine learning and deep learning. *Electronic Markets*, 31(31), pp.685–695. doi:https://doi.org/10.1007/s12525-021-00475-2.

Khan, A.N., Khan, S., Ali, A., Mansoor, M., Junaid, M.A., Jameel, A., Rustam, R., Aslam, R. and Shah, J. (2025). Artificial Intelligence in Computer Science: Evolution, Techniques, Challenges, and Multidisciplinary Applications. *Scholars Journal of Engineering and Technology*, [online] 13(04), pp.246–263. doi:https://doi.org/10.36347/sjet.2025.v13i04.006.

Khan, W., Daud, A., Khan, K., Muhammad, S. and Haq, R. (2023). Exploring the frontiers of deep learning and natural language processing: A comprehensive overview of key challenges and emerging trends. *Natural Language Processing Journal*, [online] 4(100026), p.100026. doi:https://doi.org/10.1016/j.nlp.2023.100026.

Kukhareva, P.V., Weir, C., Del Fiol, G., Aarons, G.A., Taft, T.Y., Schlechter, C.R., Reese, T.J., Curran, R.L., Nanjo, C., Borbolla, D., Staes, C.J., Morgan, K.L., Kramer, H.S., Stipelman, C.H., Shakib, J.H., Flynn, M.C. and Kawamoto, K. (2022). Evaluation in Life Cycle of Information Technology (ELICIT) framework: Supporting the innovation life cycle from business case assessment to summative evaluation. *Journal of Biomedical Informatics*, 127, p.104014. doi:https://doi.org/10.1016/j.jbi.2022.104014.

Mienye, I.D., Obaido, G., Jere, N., Mienye, E., Aruleba, K., Emmanuel, I.D. and Ogbuokiri, B. (2024). A survey of explainable artificial intelligence in healthcare: Concepts, applications, and challenges. *Informatics in Medicine Unlocked*, [online] 51, p.101587. doi:https://doi.org/10.1016/j.imu.2024.101587.

Mienye, I.D. and Swart, T.G. (2024). A Comprehensive Review of Deep Learning: Architectures, Recent Advances, and Applications. *Information*, 15(12), p.755. doi:https://doi.org/10.3390/info15120755.

Moullin, J.C., Dickson, K.S., Stadnick, N.A., Albers, B., Nilsen, P., Broder-Fingert, S., Mukasa, B. and Aarons, G.A. (2020). Ten recommendations for using implementation frameworks in research and practice. *Implementation Science Communications*, [online] 1(1). Available at: https://implementationsciencecomms.biomedcentral.com/articles/10.1186/s43058-020-00023-7.

Rasheed, K., Qayyum, A., Ghaly, M., Al-Fuqaha, A., Razi, A. and Qadir, J. (2022). Explainable, trustworthy, and ethical machine learning for healthcare: A survey. *Computers in Biology and Medicine*, 149(106043), p.106043. doi:https://doi.org/10.1016/j.compbiomed.2022.106043.

Santos, M.R., Guedes, A. and Sanchez-Gendriz, I. (2024). SHapley Additive exPlanations (SHAP) for Efficient Feature Selection in Rolling Bearing Fault Diagnosis. *Machine Learning and Knowledge Extraction*, 6(1), pp.316–341. doi:https://doi.org/10.3390/make6010016.

Smith, J.D., Li, D.H. and Rafferty, M.R. (2020). The Implementation Research Logic Model: a Method for planning, executing, reporting, and Synthesizing Implementation Projects. *Implementation Science*, [online] 15(1). Available at: https://link.springer.com/article/10.1186/s13012-020-01041-8.

Sundaram, J., Gowri, K., Devaraju, S., Gokuldev, S., Jayaprakash, S., Anandaram, H., Manivasagan, C. and Thenmozhi, M. (2023). An Exploration of Python Libraries in Machine Learning Models for Data Science. *Advances in computational intelligence and robotics book series*, pp.1–31. doi:https://doi.org/10.4018/978-1-6684-8696-2.ch001.

Vainio-Pekka, H., Mamia Agbese, Jantunen, M., Ville Vakkuri, Mikkonen, T., Rousi, R. and Pekka Abrahamsson (2023). The Role of Explainable AI in the Research Field of AI Ethics. *ACM transactions on interactive intelligent systems*, 13(4). doi:https://doi.org/10.1145/3599974.

Varajão, J., Lourenço, J.C. and Gomes, J. (2022). Models and methods for information systems project success evaluation – A review and directions for research. *Heliyon*, [online] 8(12), p.e11977. doi:https://doi.org/10.1016/j.heliyon.2022.e11977.

Vimbi, V., Shaffi, N. and Mahmud, M. (2024). Interpreting artificial intelligence models: a systematic review on the application of LIME and SHAP in Alzheimer's disease detection. *Brain informatics*, 11(1). doi:https://doi.org/10.1186/s40708-024-00222-1.

Zhang, Q. and Wang, T. (2024). Deep Learning for Exploring Landslides with Remote Sensing and Geo-Environmental Data: Frameworks, Progress, Challenges, and Opportunities. *Remote Sensing*, [online] 16(8), p.1344. doi:https://doi.org/10.3390/rs16081344.

Zhu, S., Yu, T., Xu, T., Chen, H., Schahram Dustdar, Sylvain Gigan, Deniz Gunduz, Hossain, E., Jin, Y., Lin, F., Liu, B.-N., Wan, Z., Zhang, J., Zhao, Z., Zhu, W., Chen, Z., Tariq Salim Durrani, Wang, H., Wu, J. and Zhang, T. (2023). Intelligent Computing: The Latest Advances, Challenges, and Future. *Intelligent Computing*, 2. doi:https://doi.org/10.34133/icomputing.0006.