The Socio-Technical Implications of Artificial Intelligence in Medical Image Analysis

A Research Paper submitted to the Department of Engineering and Society

Presented to the Faculty of the School of Engineering and Applied Science

University of Virginia • Charlottesville, Virginia

In Partial Fulfillment of the Requirements for the Degree

Bachelor of Science, School of Engineering

Kaihil Patel

Spring 2024

On my honor as a University Student, I have neither given nor received unauthorized aid on this assignment as defined by the Honor Guidelines for Thesis-Related Assignments

Advisor

Bryn E. Seabrook, Department of Engineering and Society

INTRODUCTION TO AI-DRIVEN HEALTHCARE

In the rapidly advancing technological landscape, machine learning and artificial intelligence (AI) emerge as pivotal forces, intricately integrating with numerous facets of human life and various sectors, driving forward the frontier of innovation and enhancing operational efficiency with unparalleled precision. One impactful use for computer vision, a specialized field of machine learning, is in the healthcare industry through medical imaging. However, there is growing concern about the potential bias of the predictions the machine learning models make (Neuroscience News, 2023). These may include biases based on protected attributes like race, sex, or age. To break down the issue of bias, a deeper analysis is taken into the data gathering and relationships between involved actors such as developers and doctors. Using the Actor-Network Theory (ANT) framework defined by Latour (1992) in *Where Are the Missing Masses? The Sociology of a Few Mundane Artifacts* will help understand the causes of implicit bias and how it impacts the use of computer vision by healthcare imaging centers in the United States.

RESEARCH METHODS AND STRUCTURE

This paper gathers sources from journals, articles, and research papers, found primarily through online databases. These sources contain information regarding the current uses of computer vision in hospitals and imaging centers, legislation and regulations on privacy, and perspectives on the future of Artificial Intelligence (AI) in the medical field. In addition, data will be gathered via an interview with the current Director of Clinical Operations at the University of Virginia Imaging Center, Amy Isakon. Interview questions were used to fill in gaps in the research conducted and obtain information that could not be otherwise concluded. These questions included but were not limited to: what are the key factors in successfully training healthcare professionals to use AI-enhanced medical imaging tools, how is AI currently being used in imaging centers, and what background knowledge do users have on AI-enabled technology? Research for the paper can be divided into sections. The first is the current approaches of automation in the healthcare industry, evaluating developments and implementations from the ANT framework perspective. Second, the relationships between relevant actants are analyzed to determine the importance and effects on other aspects of the system. Along the way, bias in the current system will be investigated in the technology itself and the bias in training and use. Finally, the consequences of bias and limitations will be discussed to provide an overview of the reality of implementation.

FOUNDATIONS OF ARTIFICIAL INTELLIGENCE IN HEALTHCARE

To accomplish this bias breakdown, it is important to understand the machine learning models in use and see what effects the training of these models has on their predictions. On a fundamental level, machine learning is a way for computers to learn and make decisions by themselves by using data, without being directly told how to do it. In the context of Computer Vision, it utilizes a common algorithmic process called a Convolutional Neural Network (CNN). CNNs are a type of artificial neural network used primarily for image recognition and processing due to their strong ability to recognize patterns in images. It attempts to mimic the human visual cortex by processing images through multiple layers that detect and recognize various features, starting from simple edges to more complex shapes (Voulodimos et al., 2018). It then uses filters to scan the image in small sections, capturing important details while gradually reducing the image size to focus on the most important parts. This process helps the CNN learn and distinguish different objects in the provided image.

If a human is tasked with identifying a unique object they have never seen before, they are likely to misidentify it. Similarly, a CNN has a higher chance of misclassification if it was not trained to identify those unique objects (Uchida et al., 2016). The lack of substantial and diverse training data is one of the largest issues with machine learning models overall. Training a model involves supplying it with labeled images and allowing it to recognize patterns and features among commonly labeled data, adjusting its parameters and improving its classification accuracy throughout. Essentially, the quality and diversity of the training data directly influence the model's performance, enabling it to generalize well to new, unseen images. Without comprehensive and representative training data, a CNN might struggle to make accurate predictions, highlighting the indispensable role of robust and varied datasets in building effective computer vision models.

Computers in the healthcare field have been increasingly prevalent since their initial introduction in the 1960s (Lipkin, 1984). Pioneers in the field sought to enhance diagnostic abilities and optimize medical analysis by teaching computers to interpret complex imaging data, ranging from X-rays to MRI scans (Sarvamangala and Kulkarni, 2022). Over time, more sophisticated algorithms and neural networks emerged as automated technology became integral to the advancement of medical diagnostics, paving the way for precision medicine and personalized treatment plans. The integration of computer vision and automation in healthcare not only streamlined the analysis of vast datasets but also significantly improved the accuracy and speed of diagnosis. However, with these advancements emerged the critical issue of bias in the models. From a machine learning perspective, model bias refers to systematic errors or unfair discrimination in the model's predictions (Mehrabi, et al., 2016). For uses in healthcare, this bias can have dire consequences. There is increased potential for misdiagnosis or delayed treatment for marginalized groups in addition to the legal and ethical implications.

ANALYZING ETHICAL DYNAMICS WITH ACTOR-NETWORK THEORY

The ethical implications of computer vision models in healthcare can be studied using Latour’s (1992) ANT. He develops a framework for exploring the dynamics and networks that shape social interactions and knowledge construction. ANT argues that society is composed of interconnected networks of both human and non-human actors, each influencing and shaping each other. He also emphasizes that relationships and networks, rather than individual entities, are crucial in understanding social processes. Regarding computer vision in the healthcare industry, there are key actor relationships between the technology (the vision models) and the users (medical professionals like radiologists).

Critics of ANT argue that it can lead to an oversimplification of complex human and societal factors by granting equal agency to non-human actors. Some scholars, like Neyrat (2018) feel that its approach might obscure important power dynamics and human responsibilities, especially in ethical and political contexts. Others like Elder-Vass (2015) criticize it for its descriptive nature, arguing that while it offers rich descriptions of specific situations, it lacks predictive or prescriptive power. Taking into consideration these criticisms, ANT is a sufficient framework to examine bias in computer vision for medical imaging.

Supporters of ANT appreciate its ability to deconstruct the human-centeredness of traditional sociological theories and its potential to illuminate the roles of technology and objects in shaping social relations. Scholars like Selbst, et al. (2019) argue that ANT provides a valuable toolkit for tracing the connections and associations between different entities, enabling a more holistic understanding of socio-technical systems. Many of these supporters relate ANT to computer algorithms as that is an exponentially growing field with substantial cause-and-effect relationships between developers, users, and the software. Supporters agree that Latour's ANT framework is a logical and rational methodology to explore the implications of algorithmic software, like computer vision, for real-world applications. The value that ANT provides is the ability to study the non-human elements of a technological system which are fundamentally vital in a computer vision model. It also divulges into the developmental process of co-production in modern advancements. In the realm of healthcare and medicine, citizens' lives are at stake and any potential risk due to technology must be mitigated.

CHALLENGES AND PROGRESS IN AI-ENHANCED MEDICAL IMAGING

While the potential of AI in medical imaging research is highly promising, several hurdles must be overcome for it to be fully integrated and utilized in clinical environments. One significant obstacle is the scarcity of high-quality, extensive, longitudinal data with outcomes. Despite dealing with the same type of disease and imaging modality, the imaging parameters and protocols can vary across clinical settings, each linked to a specific clinical scenario. The sheer number of possible clinical scenarios and the diverse tasks that images may represent are vast, presenting a challenge that might be overwhelming for any single organization using an AI algorithm. Furthermore, patient groups vary by clinic, as do clinical practices, complicating the standardization of data collection across different clinical practices. Organizing medical imaging data in a standardized manner represents a significant challenge and should itself be considered a primary field of research.

Moreover, the curation of medical imaging data faces its own set of challenges, with data curation being a crucial step that requires accurate image labeling. The exponential growth in the volume of images presents a challenge for clinicians to maintain efficiency and accuracy in processing them. Training individuals to become experts in data labeling can take years, and the difficulty in labeling a vast number of images limits the effectiveness of data curation. However, the extensive training is worth the time and effort as it will ensure the safe and proper use of AI technologies.

At the policy level, concerns about patient privacy are escalating. Strict privacy policies have historically protected patient health information, limiting the sharing of images across institutions. Recent high-profile data breaches and security attacks have led hospitals to become increasingly vigilant about security and liability, tightening security and data-sharing policies. Nonetheless, the successful implementation of AI in medical imaging necessitates access to large datasets from multiple institutions, making the secure sharing of images a challenge worth fighting for.

The Current Approach

The current network of technologies in computer vision and AI for medical imaging has seen remarkable advancements, transforming how healthcare professionals diagnose, treat, and manage various medical conditions. These technologies leverage sophisticated algorithms and deep learning models to analyze medical images such as X-rays, Magnetic Resonance Imaging (MRIs), Computed Tomography (CT) scans, and ultrasound images with high accuracy and efficiency (Esteva et al., 2021). One of the most significant advancements is the use of CNNs for image recognition and analysis. CNNs are used particularly for processing visual information and have become the standard for developing AI applications for medical imaging. They can identify patterns and features in medical images that may be subtle or invisible to the human eye. This capability is critical for early detection of diseases such as cancer, neurological disorders, and cardiovascular conditions, where early diagnosis can significantly improve patient outcomes. By analyzing vast datasets of medical images, AI models can predict the likelihood of disease progression, response to treatment, and even the risk of recurrence. This predictive power supports personalized medicine approaches, allowing healthcare providers to tailor treatments to individual patients based on their unique risk profiles and predicted disease trajectories.

The development process for AI/computer vision technology in medical imaging is multifaceted, involving a series of steps from conceptualization to clinical deployment. The first step involves clinicians, radiologists, and healthcare providers identifying a specific need or problem within medical imaging that can be addressed with or made more efficient with automated technology. From there, it enters the rest of the development cycle as discussed below.

Development Cycle: Data Collection and Labeling

The software development cycle is an iterative process that is constantly revisited and improved for as long as the software is in use. Once the problem identification and scope definition have set the goals and objectives of the AI application, the largest and most crucial step is invoked: data collection and preparation. For AI algorithms, the collection, preparation, and utilization of data are foundational processes that underscore the development and effectiveness of these technologies (Mitrofansike, 2024). This collective process, often intricate and resource-intensive, involves multiple stakeholders and a variety of data sources.

Medical imaging data for AI applications is predominantly collected from hospital archives and clinical databases. This approach has been supported by various initiatives, including the creation of platforms like The Cancer Imaging Archive (TCIA). TCIA was established to provide a central host for cancer-related imaging and data, facilitating the sharing and reuse of data for research purposes. In its first year of operation, TCIA accumulated 23 collections comprising 3.3 million images (Clark et al., 2013). The data collection effort is part of a larger initiative encouraged by the National Institutes of Health (NIH) to support secondary research through data sharing, highlighting the importance of publishing clinical and imaging data as part of fulfilling grant obligations.

For any machine learning algorithm, data is the most important aspect; it is the foundation upon which the model is built, learns, and makes predictions. The quality and quantity of the training, validation, and testing data are directly related to the accuracy of the prediction algorithm. The methodological symmetry between data quality is just as vital as the users themselves. This non-human actor should be given equal weight in the network, thereby providing a balanced view of perspectives and how bias can arise. While the current data collection methods from hospital archives and clinical databases might seem sufficient, they pose gaps in the system.

As Cath points out, AI systems utilize statistical learning techniques to identify patterns in vast datasets and make predictions based on these patterns. Consequently, if biases exist within the data, they will not only be mirrored but potentially amplified by AI algorithms in their outputs (Cath, 2018). This phenomenon complicates the ethical deployment of artificial intelligence, making it challenging to separate the biases from the algorithms themselves. In alignment with ANT, AI algorithms are not distinct from their data and should be assessed codependently. The current approach overlooks the crucial fact that biases in the data will inevitably be reflected in the AI's output, preserving the misclassification of minority groups (Buolamwini & Gebru, 2018). Data collection is not a black box and should not be overlooked. Opening the black box of data can expose and correct biases in the current approach. A global approach, involving groups of extremely diverse people must be incorporated into the dataset for the models to understand the disparities. Every human is built differently, and factors such as race, age, sex, geographical location, and socioeconomic status all play a role in how the insides of their bodies look. A vast range of demographics, stages, and abnormalities must be included to gain a comprehensive understanding of a particular disease or condition.

Once data is collected, preparation for use includes labeling the images. Image labeling relies upon human involvement to annotate imaging data with relevant information that models use to learn and make predictions. For medical images, such labeling might include the identification and marking of regions of interest, such as tumors in an MRI scan, calcifications in mammograms, or fractures in X-ray images. The labeling process requires the expertise of radiologists and medical specialists for accurate annotation and labeling, highlighting the critical features within the images that the models need to learn. A model trained on a well-labeled, diverse dataset is more likely to generalize well across different patient populations, imaging techniques, and medical facilities. It bridges the gap between raw imaging data and actionable AI insights, enabling models to learn from the expertise of medical professionals. The accuracy, diversity, and clinical relevance of these labels are paramount for creating AI tools that can effectively support diagnostic processes, ultimately contributing to better patient outcomes. Data labeling is one of the first interactions between the model and a human user and can be seen as a form of ANT translation, where human actors encode their understanding and biases into the model through labeling data. This actor relationship is critical in having human expertise to assist the models in understanding the people, and types of people, it will be identifying in the future, so extensively labeled data is better for the models to mitigate labeling bias.

Development Cycle: Bridging Data to Software

However, data must be de-identified to remove personal information, adhering to privacy regulations like the Health Insurance Portability and Accountability Act (HIPAA) (U.S. Department of Health and Human Services, 2008). Enrolling outside actants of governmental agencies and regulations possess insider influence on the system as a whole. Having well-labeled and an abundance of data contradicts the idea of data privacy which is a right by law in most countries. Recent increases in healthcare data breaches significantly impact the healthcare provider as well as the patients themselves (Seh et al., 2020). Developed countries, where healthcare data is higher in quality and abundance, incur a larger cost of data breaches. This discrepancy is why preventive measures need to be prioritized by researchers, security experts, and healthcare organizations as data is being labeled to ensure that it adheres to regulations without forfeiting data quality.

Their involvement reconfigures the network by embedding legal, ethical, and operational standards that all other actors need to adhere to, thus transforming the network's topology and altering its flows of information and power. The influence of governmental agencies and regulations extends beyond mere compliance; they act as key mediators within the network. Their input can initiate new actor relationships, dissolve existing ones, or modify the role and importance of certain actors within the network, altering the dynamics and effectiveness of the technology. On the other hand, their non-human components, such as intellectual property, subsidies, and protection laws, can catalyze innovation. Keeping governmental organizations and regulatory agencies involved in the development process amends the actor script and ensures that a system of checks and balances is maintained for data labeling. In addition, communication and addressing concerns from a developmental and technological standpoint are important in the creation and amendments to healthcare laws. Regulations cannot be made without significant input from the users themselves, and the users must adhere to regulations throughout. This symbiotic relationship between regulatory actants and the rest of the network exemplifies the ANT principle of heterogeneity, where human and non-human actors, including laws and policies, unite to form a socio-technical system. The inclusion of these actants emphasizes the importance of considering how legal and regulatory frameworks are integral to shaping technological landscapes, influencing design priorities, operational practices, and the broader societal implications of technology.

The labeled images are then delivered to the data scientists and software developers to train the models. This network delegation is a critical point in the process since it represents a shift in expertise from medical knowledge to algorithm development knowledge. Algorithm development itself is a straightforward process. Data scientists, AI researchers, and computer vision experts work collaboratively to select and refine the most appropriate machine learning models, often leveraging advanced neural network architectures such as CNNs specifically tailored for image recognition tasks. The sub-development process is iterative, focusing on model optimization to accurately interpret medical images by learning from the labeled data. Iterations include adjusting the architecture, tuning hyperparameters, and applying techniques like data augmentation and transfer learning to improve generalization and performance (Iqbal et al., 2022). The ultimate goal is to create robust, accurate models that can assist medical professionals by providing reliable interpretations of medical images to optimize diagnostic processes.

Development Cycle: Deployment and Use

Throughout this iterative sub-process, validation and testing are when medical professionals are reintroduced to the development cycle. To make a model more accurate, it must be tested by independent groups of radiologists and medical professionals. Testing validates the model's predictions against new, unseen data, and software testers and quality assurance professionals ensure the software meets all functional and performance criteria. Fine-tuning is continuously done on the model’s parameters to elevate the accuracy as much as possible before deployment.

Regulatory affairs specialists and legal advisors navigate the complex landscape of medical device regulation, preparing submissions for regulatory bodies like the Food and Drug Administration (FDA) or the European Medicines Agency (EMA). Preparation includes demonstrating the safety, efficacy, and privacy compliance of the AI application, addressing any ethical implications, and obtaining necessary certifications before widespread clinical deployment. Once deployed to imaging centers and hospitals, the primary users must be trained to use the applications. As an intersection of boundary objects, the transition to user training represents another delegation handoff between the developers and the users.

Ensuring that users are well-equipped to utilize modern technologies is the key to maximizing its effectiveness and minimizing its drawbacks. Research on specific training practices is scarce. There is little public knowledge on how, when, and who receives training for the use of AI applications in medical imaging. In addition, there is a lack of public information on the extent of their training and if it includes counseling on potential bias and their accuracies. Inadequate training can and will lead to user bias. Whether intentional or not, the users themselves pose as one of the largest sources of bias in the AI for medical imaging systems. Without fully understanding the technology, how it was created, and how it should be used, the task of using it to its full power becomes almost impossible. This situation underscores the ANT principle of translation, where the process of aligning the interests and competencies of various actors is critical for the seamless functioning of the network. Proper training ensures that users are not just operators but informed participants who understand the nuances of the technology, including its inherent biases and limitations. Currently, these technologies should be used as an assistance to the medical professional expertise since they contain limitations and biases as outlined above. Understanding the biases and interpreting them patient-by-patient is a fundamental skill that must be communicated to both users and patients.

In addition, effective integration of AI into clinical practice requires understanding not just how to operate the software but also when and why to rely on AI-assisted diagnostics. Training helps clinicians incorporate AI tools seamlessly into their workflow, enhancing productivity without compromising patient care. According to Amy Isakon, the integration of new technologies into the current workflow is one of the largest barriers to expanding AI in imaging centers (Isakon, personal communication, 2024). The field of AI in medical imaging is rapidly evolving, so constant re-training coupled with understaffing and increased demand is generating pushback on technological integration. However, improved efficiency for the short-staffed team is the factor balancing that pushback. Ongoing training helps medical professionals stay current with the latest developments, ensuring they can leverage the most current and effective tools in their practice. Physician-by-physician, there is varying knowledge about AI technology and how it is created or can be improved. Ensuring that users fully understand the potential and limitations of AI technology will allow them to stay up to date, expand AI capabilities, and have trust in their models.

Connecting to ANT's concept of mediators versus intermediaries, trained users act as mediators, actively interpreting and adapting the AI's outputs, rather than merely acting as intermediaries that pass information unfiltered. Therefore, training emerges not just as a mechanism for skill enhancement but as a pivotal factor that shapes the relationships between actors within the network, affecting the deployment, reception, and ultimate success of AI technologies in medical imaging. This interaction between the training of users and the performance of the technology exemplifies the dynamic nature of actor networks, where human capabilities and technological functionalities are intertwined, shaping the evolution and impact of the other. By understanding and compensating for the technology's biases, users can mitigate potential harms and enhance the system's reliability and effectiveness.

Consequences of Bias

One of the primary concerns is the risk of inaccurate diagnoses. When AI models are trained on datasets that are not representative of the broader patient population, they may perform poorly on images from underrepresented groups. Underrepresentation of groups in data and labeling can lead to higher rates of false positives or false negatives for these populations, directly affecting patient outcomes. For example, if a model is less accurate in detecting a particular condition in women or minority groups because it was predominantly trained on data from white males, those groups are at a higher risk of misdiagnosis. Misdiagnosis not only impacts the immediate health of the individuals concerned but also contributes to broader patterns of inequality in healthcare outcomes.

Moreover, biases in AI can exacerbate existing health disparities. Healthcare systems already struggle with disparities due to factors like socioeconomic status, geographic location, and racial or ethnic background. AI systems that are not carefully designed to account for these variables can inadvertently prioritize certain groups over others, further deepening these inequalities. For instance, an AI model trained predominantly on data from well-resourced urban hospitals may not perform as well when used in rural or resource-poor settings, potentially widening the gap in healthcare quality between these areas.

The presence of bias in AI algorithms also erodes trust in medical technology. Patients and healthcare providers must have confidence in the accuracy and fairness of AI-assisted diagnostics. When biases lead to errors or perceived inequities, it can result in a loss of trust not only in AI technologies but in the healthcare system more broadly. This is particularly concerning in contexts where AI is poised to play a significant role in diagnostic processes, as patient and practitioner trust is paramount for the successful integration of these technologies into clinical practice.

In addition, even if the bias does not come from the technology itself, it must still pass the final test of human interpretation. Training and misuse of AI technologies in medical imaging can still lead to the consequences outlined above, further harming and breaking trust between patients and the healthcare industry.

Limitations

Specific information on the data and training was limited. Due to most of the AI technology in medical imaging being developed by the private sector, intellectual property is their most important asset. It is not in their best interests to release information about the way their technology is developed to potential competitors. Furthermore, practices in training and interpretation of the results from these technologies vary by industry and location. Access to proper funding, skilled doctors, and the demographic of primary patients all play a significant role in how AI applications are used. Although one interview was completed with the Director of Imaging at the University of Virginia Health Center in which she stated AI is mostly used in MRIs, CT scans, and X-rays, this may not be entirely representative of other parts of the country or the world.

A proposal to mitigate bias and enhance the effectiveness of AI applications in medical imaging is to establish interdisciplinary research teams comprising computer scientists, healthcare professionals, lawmakers, and ethicists who can collaborate on developing more diverse and inclusive datasets that reflect the wide spectrum of patient populations. However, a limiting factor to this proposal is that including every affected actor group in a development process could not be resourcefully economical, but attempting to capture the most impactful ones is. Additionally, machine learning researchers can explore novel algorithms designed to identify and correct biases within existing models, ensuring equitable performance across different demographic groups.

CONCLUSION

The effects of bias in AI for medical imaging are multifaceted and can have profound implications on patient care, healthcare equity, and trust in medical technologies. Bias in AI algorithms can originate from various sources, including the raw data used to train these models, the data labeling for model training, and the clinical training in which they are deployed. These biases can lead to inaccurate diagnoses, reinforce existing health disparities, and undermine the potential benefits of AI in healthcare. Establishing a strong, symmetrical network with the interconnection of the various groups and actors who participate in the development of AI applications in healthcare can create a more comprehensive and understanding approach to building an effective technology. If any one party underperforms or new actors are not introduced to the process properly, implicit bias becomes more likely to infiltrate the system. The lifecycle of computer vision AI technology, especially in the healthcare field, requires intricate attention, care, and participation by a diverse group of parties all working in unison towards a common goal: efficient and improved healthcare services.

References

Buolamwini, J., Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. Proceedings of Machine Learning Research, 81(1), 77-91. http://proceedings.mlr.press/v81/buolamwini18a.html?mod=article_inline.

Cath, C. (2018). Governing artificial intelligence: Ethical, legal and technical opportunities and challenges. Philosophical Transactions Royal Society A. http://dx.doi.org/10.1098/rsta.2018.0080.

Clark, K., Vendt, B., Smith, K., Freymann, J., Kirby, J., Koppel, P., Moore, S., Phillips, S., Maffitt, D., Pringle, M., Tarbox, L., & Prior, F. (2013). The Cancer Imaging Archive (TCIA): Maintaining and Operating a Public Information Repository. Journal of Digital Imaging, 26, 1045-1057. https://doi.org/10.1007/s10278-013-9622-7.

Elder-Vass, D. (2014). Disassembling actor-network theory. Loughborough University, https://hdl.handle.net/2134/17424

Esteva, A., Chou, K., Yeung, S., Naik, N., Madani, A., Mottaghi, A., Liu, Y., Topol, E., Dean, J., & Socher, R. (2021). Deep learning-enabled medical computer vision. NPJ Digital Medicine, 4, https://doi.org/10.1038/s41746-020-00376-2.

Iqbal, S., Adnan N. Qureshi, A.U., Li, J., Mahmood, T. (2022). Improving the Robustness and Quality of Biomedical CNN Models through Adaptive Hyperparameter Tuning. Applied Sciences, 12(22), https://doi.org/10.3390/app122211870

Latour, B. (1992). Where Are the Missing Masses? The Sociology of a Few Mundane Artifacts. In W. E. Bijker & J. Law (ed.), Shaping Technology / Building Society: Studies in Sociotechnical Change, 152-179.

Lipkin M. (1984). Historical Background on the Origin of Computer Medicine. Proceedings of the Annual Symposium on Computer Application in Medical Care, 987–990.

Mehrabi, Ninareh; Morstatter, Fred; Saxena, Nripsuta; Lerman, Kristina; Galstyan, Aram. (July 2021). A Survey on Bias and Fairness in Machine Learning. ACM Computing Surveys, 54(6), 115:1-115:35. https://doi.org/10.1145/3457607.

Mitrofansiky, Kosta. (2024, February 28). Medical Imaging Software Development: a Detailed Guide. Intellisoft. https://intellisoft.io/medical-imaging-software-development-a-detailed-guide/.

Neuroscience News. (2023, April 26). Unveiling the Hidden Biases in Medical AI: Paving the Way for Fairer and More Accurate Imaging Diagnoses. Neuroscience News, https://neurosciencenews.com/ai-medical-imaging 23117/#:~:text=One%20of%20the%20main%20sources

Neyrat, F., D. S. (2019). The Unconstructable Earth: An Ecology of Separation. Fordham University Press, https://doi.org/10.2307/j.ctv8jnzp4

Obermeyer, Z., Powers, B., Vogeli, C., Mullainathan, S. (2019). Dissecting racial bias in an algorithm used to manage the health of populations. Science, 366(6464), 447-453. https://doi.org/10.1126/science.aax2342.

Pinto-Coelho, L. (2023). How Artificial Intelligence Is Shaping Medical Imaging Technology: A Survey of Innovations and Applications. National Library of Medicine, 10(12), https://doi.org/10.3390/bioengineering10121435.

Sarvamangala, D. R., & Kulkarni, R. V. (2022). Convolutional Neural Networks in Medical Image Understanding: a Survey. Evolutionary Intelligence, 15(1), 1–22, https://doi.org/10.1007/s12065-020-00540-3

Seh, A.H., Zarour, M., Alenezi, M., Sarkar, A.K., Agrawal, A., Kumar, R., Ahmad, K.R. (2020) Healthcare Data Breaches: Insights and Implications. Healthcare, 8(2), 133. https://doi.org/10.3390/healthcare8020133.

Selbst, Andrew D.; Boyd, Danah; Friedler, Sorelle A.; Venkatasubramanian, Suresh; Vertesti, Janet. (2019). Fairness and Abstraction in Sociotechnical Systems. Proceedings of the Conference on Fairness, Accountability, and Transparency, Association for Computing Machinery, 59–68, https://doi.org/10.1145/3287560.3287598

Uchida, Seiichi; Shota Ide; Brian Kenji Iwana; Anna Zhu. (2016). A Further Step to Perfect Accuracy by Training CNN with Larger Data. 2016 15th International Conference on Frontiers in Handwriting Recognition (ICFHR), 405–10. https://doi.org/10.1109/ICFHR.2016.0082.

U.S. Department of Health and Human Services. (2008, May 7). Summary of the HIPAA Privacy Rule., https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html.

Voulodimos, Athanasios; Doulamis, Nikolaos; Doulamis, Anastasios; Protopapadakis, Eftychios. (2018). Deep Learning for Computer Vision: A Brief Review. Computational Intelligence and Neuroscience, 1-13. https://doi.org/10.1155/2018/7068349

Appendix: Questions Aked to UVA Imaging Director Amy Isakon

1. How do AI and ML currently enhance medical imaging processes in healthcare, and what are the most significant improvements you have observed?
2. Can you discuss the challenges of integrating AI/ML technologies into existing healthcare workflows for medical imaging?
3. How do you see the role of AI/ML in medical imaging evolving in the next 5-10 years?
4. In your opinion, what are the key factors in successfully training healthcare professionals to use AI-enhanced medical imaging tools?