Geometric Heuristics Enhance POCUS AI for Pneumothorax

Viekash Vinoth Kumar

CMU-RI-TR-24-11

Apr 24, 2024

The Robotics Institute
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA

Thesis Committee:
Dr. John Galeotti *(chair)*
Dr. Deva Kannan Ramanan
Mr. Nishant Thumbavanam Arun

Submitted in partial fulfillment of the requirements for the degree of Master of Science in Robotics.

To my family, friends, and mentors.

Abstract

Point-of-care ultrasound (POCUS) represents a significant advancement in emergency and critical care medicine, offering real-time imaging capabilities that are essential for rapid diagnostic and therapeutic decisions. Despite its advantages, the effective interpretation of POCUS images demands a high degree of expertise, making the process susceptible to variability and potential diagnostic inaccuracies. POCUS can also be challenging for artificial intelligence (AI) to interpret, and existing commercial POCUS AI systems required large amounts of labeled training data. We trained POCUS AI models using limited training data, for which we improved diagnostic accuracy by directly integrating geometric heuristics into the architecture and training of our POCUS AI models. These heuristics were derived by humans from the expert knowledge of clinicians.

Our clinical focus was diagnosing pneumothorax, where rapid and accurate detection is crucial. Our AI-enhancing methodology centered on collecting and incorporating geometric heuristics, i.e. intuitive rules and patterns used by clinicians in image interpretation. These heuristics include the observation of pleural line sliding, its relative movement against the intercostal muscle, and the specific positioning of the pleural line, among others. We represented these heuristics using semantic-segmentation label images and optical flow images, and we also cropped the images based on the semantic segmentation. We developed two distinct methods for embedding these additional heuristic images into the AI models: one through adding new channels to the input data and another by integrating them as distinct inputs that the AI model later fuses into a common embedding space with the original grayscale image data.

The strategies of cropping regions of interest (ROI) and utilizing segmentation maps significantly outperformed the baseline models, underscoring the importance of directing the AI's focus to crucial image regions. Surprisingly, optical flow maps did not enhance model performance, highlighting the nuanced nature of computing and/or integrating motion-related heuristics for ultrasound. Overall, the multi-channel input approach proved slightly more effective than the fused embedding space, though both methods showed promise in improving diagnostic accuracy.

This study opens up avenues for exploring additional heuristic combinations and refining model architectures to further enhance the performance

of AI in medical imaging, especially in applications such as POCUS video for which both clinicians and AI often struggle to discern nondescript anatomy and motion.

Acknowledgments

This thesis stands as a milestone in my academic journey, a journey made possible through the support, wisdom, and encouragement of many. Foremost, I extend my deepest gratitude to my advisor, Professor John Galeotti, whose guidance has been nothing short of transformational. Throughout my master's program and the development of this work, Professor Galeotti has been a beacon of knowledge, providing me with unparalleled opportunities and sage advice that have profoundly shaped my academic and personal growth. His mentorship has been a cornerstone of my experience, for which I am eternally grateful.

I am also indebted to my lab colleagues Nishanth Thumbavanam Arun, Ananya Bal, and Gautam Rajendrakumar Gare. Their camaraderie, insightful discussions, and unwavering support have been instrumental in navigating the challenges of research. The environment of mutual aid and feedback in our lab has been a fertile ground for my development as a researcher.

Special thanks are due to my committee member, Professor Deva Kannan Ramanan. The courses I have taken under his tutelage at Carnegie Mellon University (CMU) have enriched my knowledge base and have been integral to my research's success. Professor Ramanan's contributions to my academic journey have left an indelible mark on my scholarly pursuits.

I extend my heartfelt appreciation to our collaborators at the University of Pittsburgh Medical Center (UPMC). Their expertise and collaborative spirit have significantly enhanced the scope and impact of our research.

My time at CMU has been an unforgettable chapter of my life, made all the more memorable by the friendships and connections I have forged. I am thankful for every moment shared with my friends and peers at CMU, whose presence made my journey not only educational but also enjoyable and enriching.

In closing, I recognize that this thesis is not merely a product of my efforts but a testament to the collective support, encouragement, and inspiration I have received from the aforementioned individuals and many others. To everyone who has played a part in my journey, your impact extends beyond the pages of this work, and for that, I am profoundly thankful.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 Motivation

The advent of Point-of-Care Ultrasound (POCUS) has marked a significant advancement in the field of emergency and critical care, offering a vital tool for rapid and effective decision-making [47]. Despite its transformative impact, the effective utilization of POCUS is heavily reliant on the operator's expertise, highlighting a crucial challenge in the field: the variability in diagnostic accuracy across different levels of clinical skill [70]. This variability underscores the need for innovative approaches to mitigate the reliance on individual expertise, aiming to democratize the benefits of POCUS across various clinical settings [34].

The versatility and non-invasiveness of POCUS make it an indispensable tool in diagnosing and managing critical conditions [36]. However, the efficacy of POCUS hinges on the clinician's ability to accurately interpret complex ultrasound images, a skill that requires extensive training and experience [33]. The inherent limitations of human interpretation, including susceptibility to fatigue and cognitive biases, further underscore the potential of artificial intelligence (AI) to augment diagnostic accuracy in POCUS applications [40].

Recent advancements in AI, particularly deep learning, have shown promise in enhancing medical imaging analysis, offering a pathway to augment human capabilities

Figure 1.1: Integration of AI in augmenting POCUS diagnostic accuracy.

and reduce variability in image interpretation [41], [8].

Our interdisciplinary approach includes both computational techniques and the clinical context of POCUS usage [7], [63]. This encompasses an appreciation of ultrasound physics, the variability of human anatomy, and the spectrum of pathologies influencing image interpretation [25], [45]. Our research is driven by the potential to significantly improve patient outcomes through enhanced diagnostic reliability and timeliness, particularly in emergency settings where rapid decision-making is critical [27], [4].

Furthermore, by enhancing POCUS with AI, we envision expanding its diagnostic benefits to a wider range of clinical settings, including those with limited access to specialized radiological expertise [80], [22]. This aligns with the broader goal of making advanced diagnostic capabilities universally accessible, ensuring equitable healthcare delivery [58], [55].

Our research is propelled by the ambition to leverage AI in unlocking the full diagnostic potential of POCUS, merging technological innovation with clinical expertise [66], [32]. By incorporating geometric heuristics into AI models, we aspire to develop systems that not only elevate diagnostic accuracy but also encapsulate the nuanced understanding characteristic of experienced clinicians, redefining patient care by broadening the accessibility and reliability of ultrasound interpretation across all levels of clinical practice [54], [28].

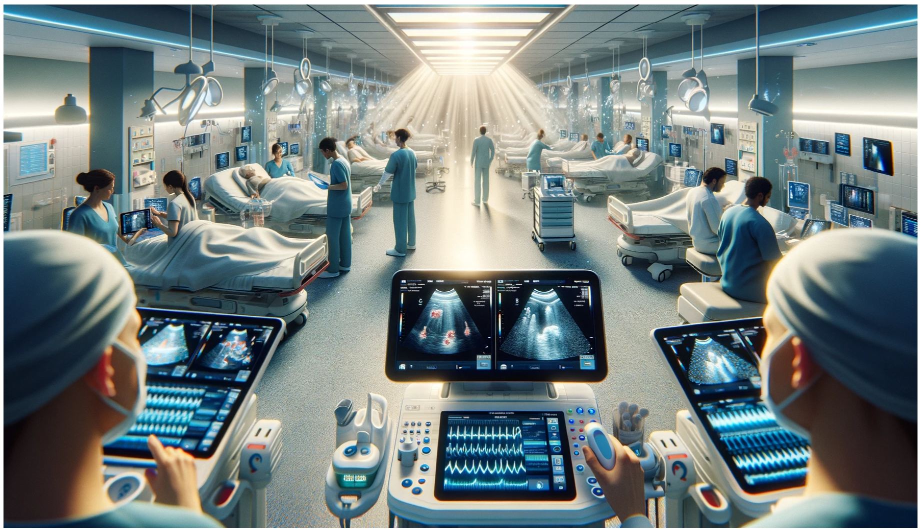

Figure 1.2: AI-augmented POCUS in various clinical settings.

1.2 Contributions

This thesis progresses in AI-driven medical diagnostics using Point-of-Care Ultrasound (POCUS) through two primary novel contributions. The first is the visualizing of critical geometric heuristics derived from clinicians. These heuristics capture essential diagnostic patterns and shapes that clinicians intuitively use during POCUS interpretation. The second major contribution is the development of a method to effectively communicate these identified heuristics to deep learning models.

Additionally, the research includes secondary contributions that further the understanding and effectiveness of heuristic integration in AI models. A comprehensive study on the effect of noise within the imaging data has been conducted to determine how it impacts the model's ability to utilize the geometric heuristics, aiming to refine the model's resilience and accuracy under real-world conditions. Moreover, rigorous testing has been performed to quantify the impact of the identified heuristics on the model's performance. These tests validate the practical significance of integrating geometric heuristics into POCUS AI models, emphasizing their value in enhancing diagnostic precision.

Chapter 2

Background and Related Works

2.1 Point-of-Care Ultrasound (POCUS)

2.1.1 Introduction to POCUS and AI

Point-of-Care Ultrasound (POCUS) has become an essential diagnostic tool in emergency and critical care, offering rapid, real-time imaging to guide clinical decisions. Despite its advantages, the accuracy of POCUS is heavily dependent on the operator's skill, presenting a challenge for its broader implementation [18, 62]. The integration of Artificial Intelligence (AI) promises to mitigate these challenges by enhancing diagnostic accuracy and making POCUS accessible to a wider range of healthcare professionals [10, 26].

2.1.2 Diagnostic Accuracy and AI Integration

Recent studies underscore the potential of AI in improving the diagnostic accuracy of POCUS. Specifically, AI-assisted assessments of left ventricular ejection fraction (LVEF) have shown high sensitivity and specificity, even when performed by less experienced operators [20, 67]. These findings highlight AI's role in standardizing POCUS interpretations and expanding its utility across diverse clinical settings.

Figure 2.1: Showcasing the transformative potential of artificial intelligence in enhancing medical imaging interpretations [35].

2.1.3 AI-Enhanced POCUS in Low-Resource Settings

In low-resource settings (LRS), AI-enhanced POCUS could significantly impact healthcare delivery by providing accurate diagnostic capabilities where access to trained specialists is limited. Research demonstrates that minimal training is required for non-experts to achieve reliable imaging results with AI-assisted POCUS, making it a valuable tool for disease diagnosis and management in underserved areas [3, 43].

2.1.4 Advancements in AI-Enhanced POCUS

Recent studies have broadened the applications of AI-enhanced POCUS, demonstrating its potential beyond cardiac assessments to include respiratory, abdominal, and vascular imaging. The integration of AI has been instrumental in improving the diagnostic accuracy for conditions such as pneumothorax, pulmonary edema, and acute respiratory distress syndrome (ARDS), where AI-POCUS systems have matched or surpassed the performance of conventional imaging modalities like chest X-ray and CT scans. Moreover, the application of AI in obstetric POCUS for fetal health monitoring and gestational age estimation has underscored the versatility of AI-POCUS across different medical specialties [6, 44].

2.1.5 Emergency and Critical Care Applications

In emergency and critical care settings, the value of AI-enhanced POCUS is particularly pronounced, offering rapid and accurate diagnostics crucial for patient management. Research highlights its utility in emergency departments for conditions requiring immediate intervention, such as cardiac tamponade and abdominal emergencies. AI assistance in POCUS can mitigate the risk of diagnostic errors, especially under the high-pressure conditions typical of emergency care, thereby enhancing patient outcomes [15, 82].

Figure 2.2: AI-enhanced Point-of-Care Ultrasound (POCUS) in Emergency and Critical Care Settings, showcasing its pivotal role in rapid and accurate diagnostics crucial for patient management [30].

2.1.6 Training and Implementation Challenges

Despite its promise, the widespread implementation of AI-enhanced POCUS encounters significant challenges, notably in training healthcare professionals to effectively use these systems. The diversity in AI-POCUS interfaces and the intricacies of AI-assisted diagnostic interpretations necessitate robust training programs. Furthermore, seamlessly integrating AI-POCUS into clinical workflows demands adaptation and compatibility with existing medical practices and infrastructure [50, 61].

2.1.7 Ethical and Regulatory Considerations

The advancement of AI applications in healthcare brings forth critical ethical and regulatory considerations. Issues related to patient privacy, data security, and the accountability of AI-assisted diagnostic decisions are of paramount concern. Ensuring that AI-POCUS systems are developed and utilized in compliance with stringent ethical standards and regulatory guidelines is essential to maintaining patient safety and trust in healthcare technology [16, 65].

2.1.8 Future Directions in Research

Looking forward, research in AI-enhanced POCUS will likely focus on exploring new diagnostic applications, enhancing user experience, and overcoming current technological limitations. Progress in machine learning algorithms and imaging technology promises more sophisticated AI-POCUS systems capable of providing comprehensive diagnostic insights. Collaborative efforts among clinicians, researchers, and technology developers are crucial for realizing the full potential of AI in transforming POCUS applications [71, 79].

In conclusion, the integration of AI into POCUS marks a significant advancement in medical imaging, with the potential to improve diagnostic accuracy, expand application scopes, and enhance patient care across a wide array of medical disciplines. The ongoing evolution of this field is set to play a critical role in the future landscape of healthcare delivery.

2.1.9 Introduction to Ultrasound in Clinical Decision-Making

Ultrasound imaging is emerging as a new cornerstone in point-of-care medical diagnostics, offering real-time visualization of internal organs and structures. Its non-invasive nature and versatility make it an indispensable tool across various medical fields. Despite its advantages, the interpretation of ultrasound images requires a high level of expertise, as it is highly operator-dependent and subject to inter- and intra-observer variability [14, 46].

2.2 Clinician Heuristics in Ultrasound Interpretation

2.2.1 Ultrasound Heuristics

Clinician heuristics, defined as cognitive shortcuts or rules of thumb that assist in decision-making, play a crucial role in the interpretation of ultrasound images. These heuristics, derived from experience, help in navigating complex diagnostic environments by simplifying the processing of visual information and facilitating rapid decision-making [76, 79].

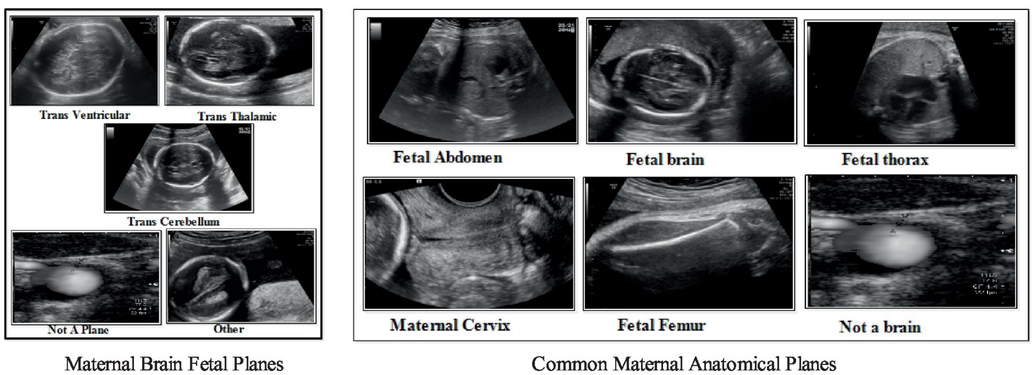

Figure 2.3: Insightful Ultrasound Heuristics: Clinically Derived Rules and Patterns Guiding Interpretation for Accurate Diagnosis [68].

Incorporating Heuristics into AI Models

Recent advancements in artificial intelligence (AI) and machine learning (ML) have shown potential in augmenting the diagnostic process, including ultrasound interpretation. Integrating clinician-derived heuristics into AI models aims to bridge the gap between human expertise and machine precision, enhancing the model's interpretability and accuracy [71, 76]. One study demonstrated the effectiveness of heuristic corrections in improving the segmentation of cardiac structures in ultrasound images, showcasing the potential of combining human expertise with AI capabilities [9].

Challenges and Opportunities

While the integration of heuristics into AI models presents a promising avenue, it also introduces challenges, including the codification of expert knowledge into a form that can be understood by AI systems and ensuring that the AI model remains adaptable to new or unseen cases [24, 60].

2.2.2 Nudges in Clinical Decision-Making

Beyond heuristics, the concept of 'nudges'—subtle interventions designed to influence behavior without restricting choices—has been explored in healthcare to guide clinicians towards improved decision-making. Nudges have been employed in various healthcare settings to encourage the adoption of evidence-based practices, adherence to guidelines, and to improve the efficiency and quality of care [29, 77].

Effectiveness of Nudges

A systematic review of clinician-directed nudges revealed their potential in positively influencing clinical decision-making, emphasizing the importance of choice architecture in healthcare settings. Nudges that alter information presentation, decision structures, and provide decision aids can guide clinicians towards better clinical decisions, thereby indirectly influencing the utilization and interpretation of diagnostic tools like ultrasound [29, 77].

The integration of clinician heuristics and the application of nudges in clinical decision-making represent two complementary approaches to enhancing the diagnostic process in ultrasound imaging. These strategies underscore the importance of leveraging both human expertise and behavioral insights to support clinical decision-making, paving the way for more accurate, efficient, and patient-centered care.

2.3 Ultrasound Video Classification

The advancement of machine learning (ML) and deep learning (DL) techniques has significantly impacted the field of medical imaging. This section delves into the evolution and current state of these technologies, focusing on their application in ultrasound video classification [2, 52].

2.3.1 Machine Learning in Ultrasound Imaging

Machine learning, with its ability to learn from data, identify patterns, and make decisions with minimal human intervention, has found extensive applications across various fields, including ultrasound imaging [2, 52]. The use of ML in ultrasound imaging is not new; however, recent years have witnessed a substantial growth in interest and application of these techniques. This surge is attributed to ML's ability to address some of the inherent challenges in ultrasound imaging, such as the interpretation of complex patterns and the need for automated, reliable diagnostics [2, 52].

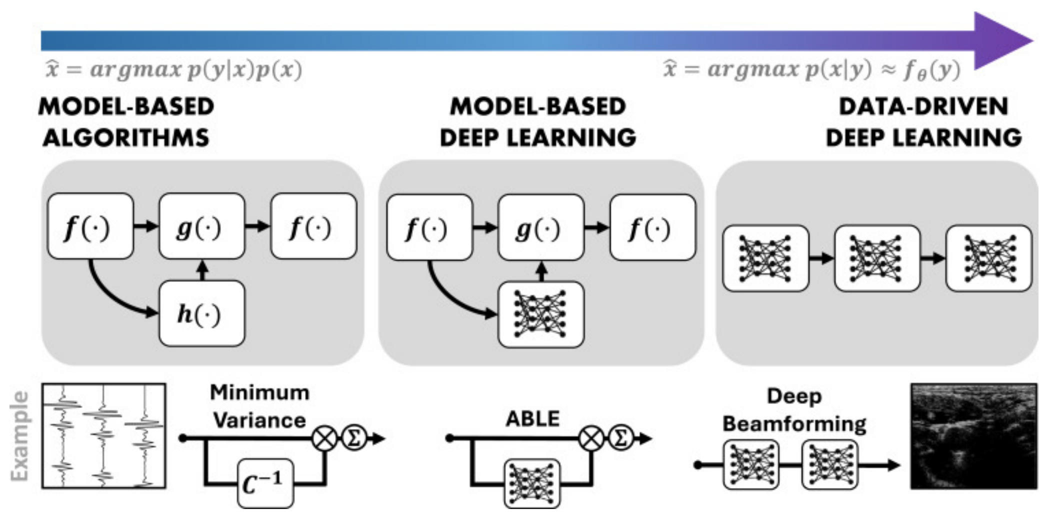

Figure 2.4: Ultrasound ML Classification: From Interpreting Complex Patterns to Ensuring Automated, Reliable Diagnostics [52].

2.3.2 Deep Learning for Medical Diagnostic Videos

Deep learning, a subset of ML, has shown remarkable success in the automated analysis of medical diagnostic videos, including ultrasound and endoscopy. By automating tasks such as classification, detection, and segmentation, DL models enhance the efficiency and accuracy of diagnoses. Convolutional Neural Networks (CNNs) and

Long Short-Term Memory (LSTM) networks are among the most commonly used models for analyzing medical videos, demonstrating significant success in extracting meaningful information from complex datasets [13, 51].

Applications in Ultrasound Video Classification

The classification of ultrasound videos using DL has particularly benefited from the capabilities of CNNs and LSTM networks. These models have been effectively applied to a range of tasks, from segmenting specific anatomical structures to classifying different types of medical conditions within ultrasound videos.

Challenges and Future Directions

Despite the successes, the application of DL in ultrasound video classification faces several challenges. These include the labeling and preprocessing of medical videos, class imbalance, and time complexity. Furthermore, integrating expert knowledge into DL models remains an open challenge, emphasizing the need for collaborative research efforts with domain experts to enhance diagnostic processes [17].

The integration of machine learning and deep learning in ultrasound imaging represents a promising frontier in medical diagnostics. As technology advances, further collaboration between computational scientists and medical professionals is essential to fully harness the potential of these models, paving the way for more accurate, efficient, and automated diagnostic tools.

Chapter 3

Experimental Setup and Data

This chapter outlines the experimental framework and datasets employed in these investigations. The study leverages a comprehensive video dataset augmented with derived heuristics datasets, including segmentation maps, Pleura ROI crops, and optical flow maps. The model selection and computational resources necessary for the research are detailed, providing a foundation for the experimental analysis.

Figure 3.1: Ultrasound Video Frame of the Lung: Visualizing Pulmonary Structures and Pathologies in Real Time.

3.1 Video Dataset

The primary dataset comprises 107 POCUS video clips sourced from a diverse patient demographic presenting with various conditions. This is a very small dataset for deep learning, but is relatively large in comparison to most publicly available standardized lung-ultrasound datasets. (There are larger lung-ultrasound datasets, but they are are diversely aggregated and lack a common protocol, making them challenging to use.) The allocation of 64 video clips for training includes only 30 positive examples of pneumothorax. An additional 43 video clips were used for testing.

At the Brooke Army Medical Center (BAMC), each video segment was acquired utilizing a Butterfly Ultrasound device equipped with a high-resolution probe. This method enabled the detailed capture of thoracic lung ultrasound scans, specifically targeted at pivotal regions for the identification of medical conditions such as pneumothorax. Prior to inclusion in the dataset, each video was anonymized to protect patient privacy, with all identifiable patient information removed.

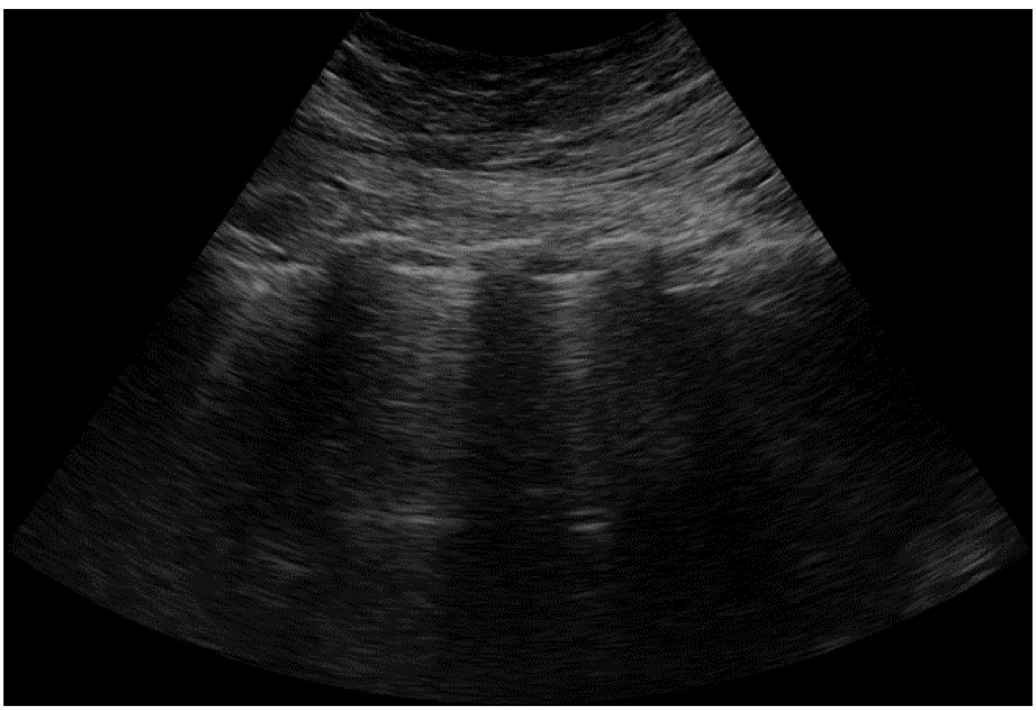

Figure 3.2: Preprocessing Step 1: Cropping the Flank Sides of the Ultrasound Image to Remove Scale and Markers, Enhancing Clarity for Subsequent Analysis.

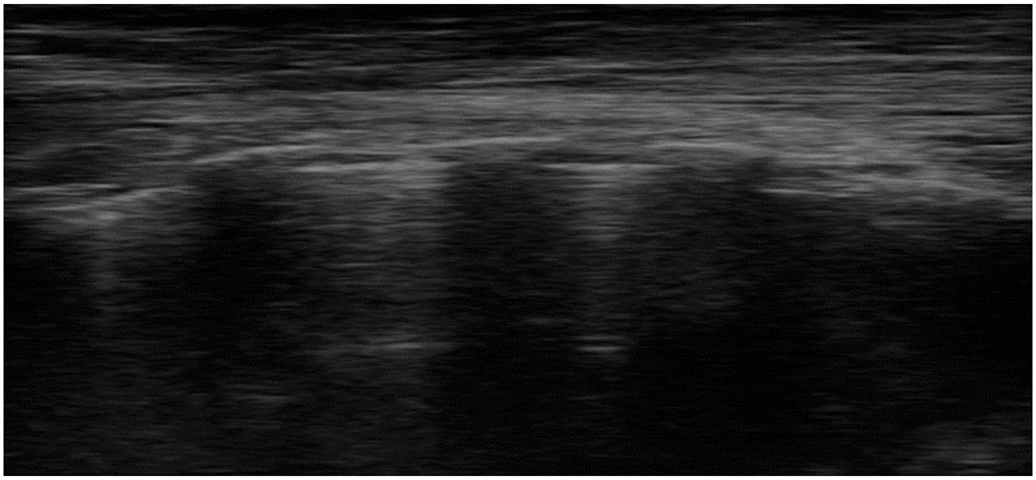

Figure 3.3: Processing Step 2: Rectilinear Transformation with Columns Representing A-scans, Transforming Ultrasound Data for Enhanced Visualization and Analysis.

3.2 Heuristics Dataset

The heuristics dataset is designed to encapsulate clinician expertise in the form of geometric patterns and landmarks critical for POCUS image interpretation. This dataset is subdivided into three categories: segmentation maps, Pleura ROI (Region of Interest) crops, and optical flow maps. Each category targets specific aspects of the ultrasound images, providing information on clinically relevant features directly to the AI model, so that the model doesn't have to independently learn to interpret some of these complexities of these medical images.

3.2.1 Segmentation Maps

Segmentation maps are generated for each video frame, delineating key anatomical structures such as the pleural line, intercostal muscles, and ribs.

3.2.2 Pleura ROI Crops

Focusing on the pleura-intercostal region, Pleura ROI crops are extracted from the ultrasound frames. This heuristic preserves only the image region around the pleural line and adjacent areas, which are critical for diagnosing pneumothorax; all other

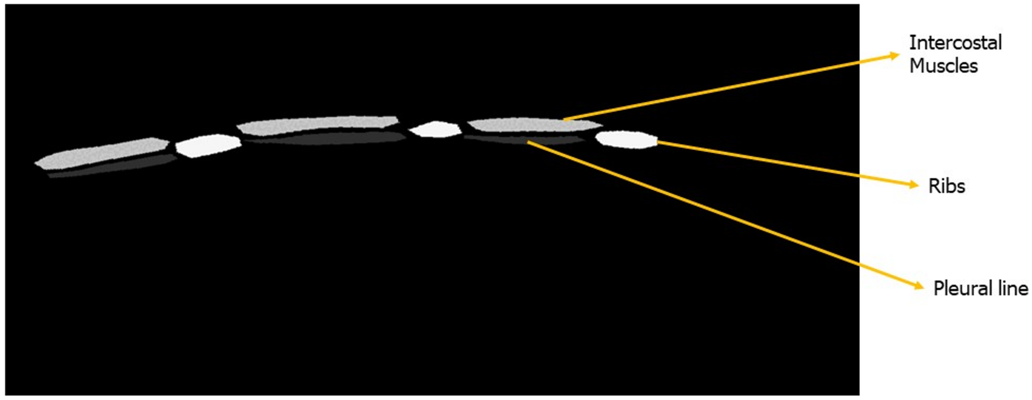

Figure 3.4: Segmentation map showing labeled structures that facilitated AI interpretation of lung ultrasound: 1) Rib, 2) Pleura Line, 3) Intercostal Muscle.

image content is discarded to avoid accidental AI over fitting to irrelevant image content. The crops are sized to include the pleural line and enough surrounding tissue to allow the model to assess the presence of pleural sliding and other diagnostic indicators.

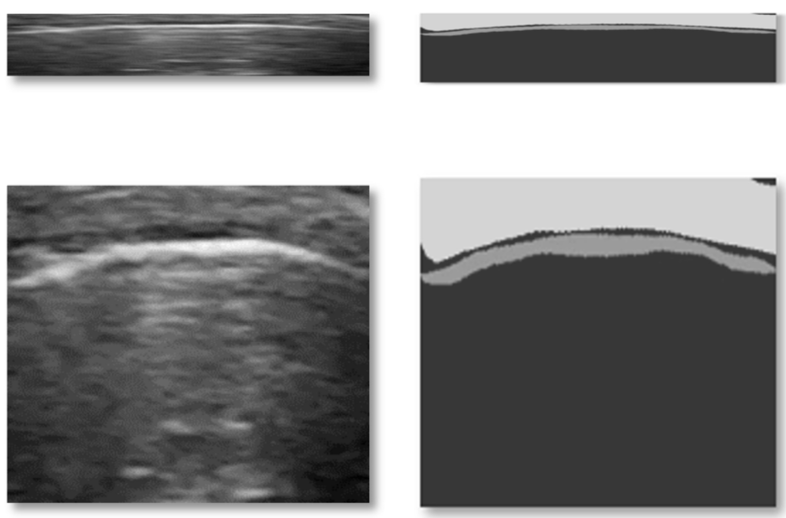

Figure 3.5: Region of Interest Cropped from Ultrasound Frame (Left) and Segmentation Map (Right), Highlighting Labeled Structures for Detailed Analysis.

3.2.3 Optical Flow Maps

Optical flow maps represent the motion between consecutive frames of the ultrasound videos, capturing the dynamic nature of the anatomical structures under observation. Optical flow maps are generated in both the X and Y directions, to represent the complete motion patterns. When used in conjunction with static segmentation maps, both spatial and temporal information is represented in the heuristics.

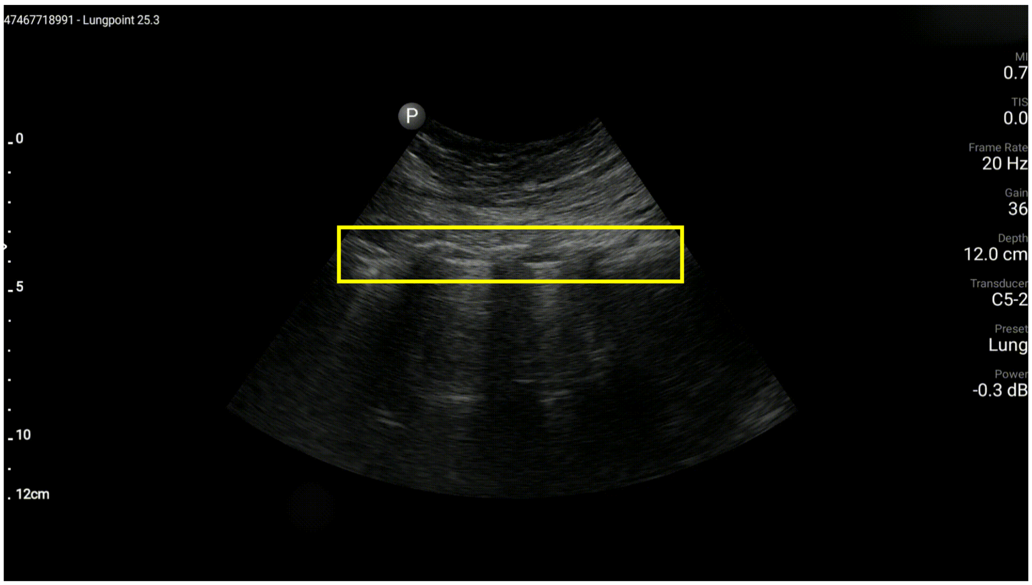

Figure 3.6: Region of Interest Containing the Sliding Motion of the Pleura.

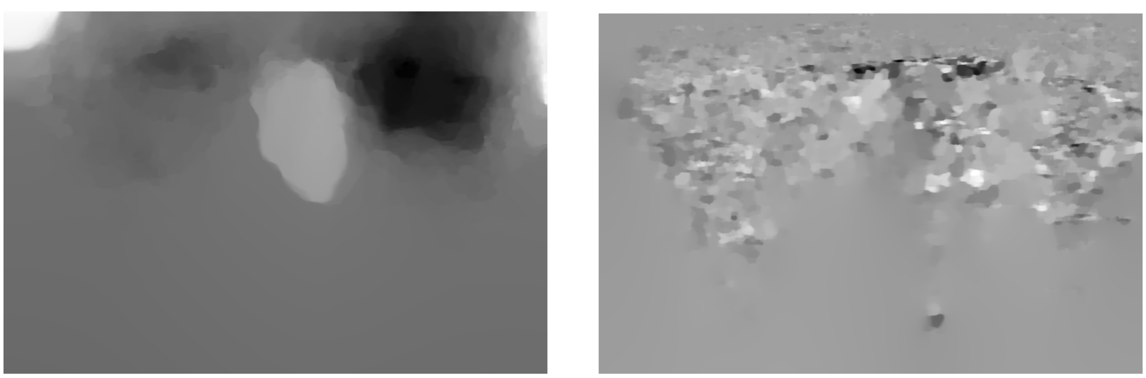

Figure 3.7: Optical Flow Maps: Left - X Direction, Right - Y Direction, Facilitating Analysis of Motion in Ultrasound Imaging.

3.3 Model Selection

The experimental framework utilizes a dual-method approach for integrating the heuristics into the AI model. Two architectures, ResNet-50 with Temporal Shift Module (TSM) and MobileNet V2 with TSM, were selected based on their proven efficacy in image classification tasks and their ability to handle temporal data. ResNet-50 with TSM is employed for its deep learning capabilities and robust feature extraction, while MobileNet V2 with TSM offers a lightweight alternative suitable for real-time applications on portable devices [5, 38].

Implementation Details: The learning-based components of our study are implemented using PyTorch, a leading deep learning framework known for its flexibility, efficiency, and support for dynamic computational graphs. PyTorch facilitates the rapid prototyping of deep learning models, allowing for direct manipulation of model architecture and training procedures.

Data augmentation techniques were tailored to mimic the variability and challenges inherent in clinical ultrasound imaging. For instance, rotations and translations simulate the different orientations and positions in which ultrasound scans may be conducted in a hectic emergency room setting, while scaling adjustments account for the variability in organ sizes among different patients.

Chapter 4

Selection of Heuristics

4.1 Introduction

The domain of artificial intelligence (AI) and machine learning (ML) has seen remarkable progress in recent years in the field of medical imaging analysis [56, 78]. Heuristics, which are rule-based approaches derived from human expertise, can be used to guide AI models to achieve higher accuracy and interpretability, especially in complex tasks like the classification of medical images [1, 69]. This chapter focuses on the selection and application of heuristics in enhancing the performance of POCUS AI models for the particular application of diagnosing pneumothorax.

4.2 What Are Heuristics?

Heuristics, in the context of artificial intelligence and machine learning, refer to simplified rules or strategies derived from human expertise intended to solve complex problems efficiently [39]. These rules, especially geometric heuristics in medical imaging, aid in narrowing down the focus areas for AI models, facilitating better interpretation of images by highlighting crucial patterns and relationships [48].

4.3 How to Choose Heuristics

Selecting appropriate heuristics is a nuanced process that involves evaluating potential rules based on several criteria, including relevance to the task, computational efficiency, and impact on model performance [11]. This section outlines the methodology employed for heuristic selection, which encompasses empirical analysis and expert consultation. It also addresses the challenges inherent in this selection process, such as ensuring the balance between complexity and performance and achieving generalizability across different datasets [53].

4.3.1 Criteria for Selection

The selection of heuristics is guided by several criteria [73], including:

1. **Relevance**: The heuristic must be directly applicable to the diagnostic problem at hand, capable of highlighting key features or patterns that are indicative of the condition being diagnosed [37].
2. **Computational Efficiency**: Selected heuristics should be computationally efficient to apply within the AI model without significantly increasing processing time, especially for edge deployment and/or real-time video applications [81].
3. **Interpretability**: The heuristic should contribute to the interpretability of the AI model's decisions, allowing clinicians to understand the basis of the AI's diagnosis and trust its accuracy [21].
4. **Impact on Model Performance**: The primary goal of integrating heuristics is to enhance the model's diagnostic accuracy [42].

4.3.2 Collaboration with Clinicians

Clinicians can provide insights into their diagnostic reasoning process, highlighting key visual cues, patterns, and geometric relationships that they rely on when interpreting POCUS images [57]. For example, clinicians might highlight the importance of the pleural line's visibility and its dynamic behavior (e.g., lung sliding) as crucial indicators of pneumothorax, for which a heuristic can be created.

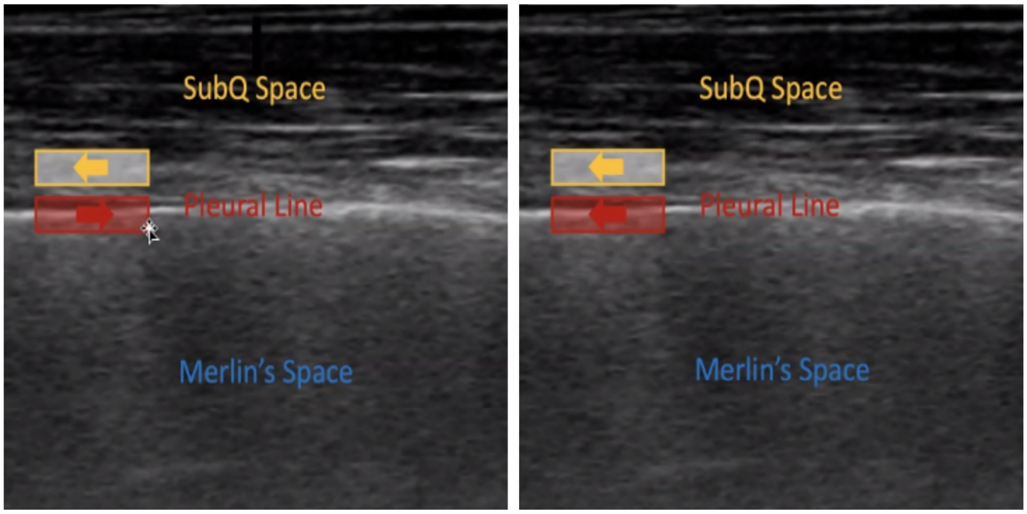

Healthy Patient   Pneumothorax Patient

Figure 4.1: Comparison Between Ultrasound Images of a Patient with Pneumothorax and a Healthy Patient.

4.4 Selecting Heuristics for the Problem Under Study: Pneumothorax Classification

Pneumothorax is an emergency situation characterized by the presence of (often actively accumulating) air in the pleural space just outside the lungs, making it increasingly difficult for the lungs to expand when breathing. This section explores key heuristics significant in pneumothorax detection in ultrasound, detailing their clinical importance and implementation in AI modeling.

Collection of Geometric Heuristics

Clinicians utilize several key heuristics for pneumothorax diagnosis in lung ultrasound images. These include: observation of pleural line sliding, comparison of pleural line movement relative to intercostal muscles, identification of pleural line positioning between ribs, examination of the pleura-intercostal region, and evaluation of sliding motion. These heuristics and their visualization methods were systematically categorized as shown in Table 4.1.

Representation of Heuristics

Effective presentation of these heuristics to AI is crucial for their integration into the AI model. For pleural line sliding, sliding relative to intercostal muscles, and pleural line positioning between ribs, segmentation image maps were used as additional inputs to the AI model. Experts created segmentation maps for each ultrasound image frame, for each of the Pleural line, Intercostal muscle, and Ribs. For heuristics focusing on the pleura-intercostal region, a ROI cropping technique was employed, involving direct cropping around the pleura from video frames; the exact ROI crop was based on the segmentation maps. For the assessment of sliding motion, Optical Flow Maps in both X and Y directions were computed using standard techniques [75].

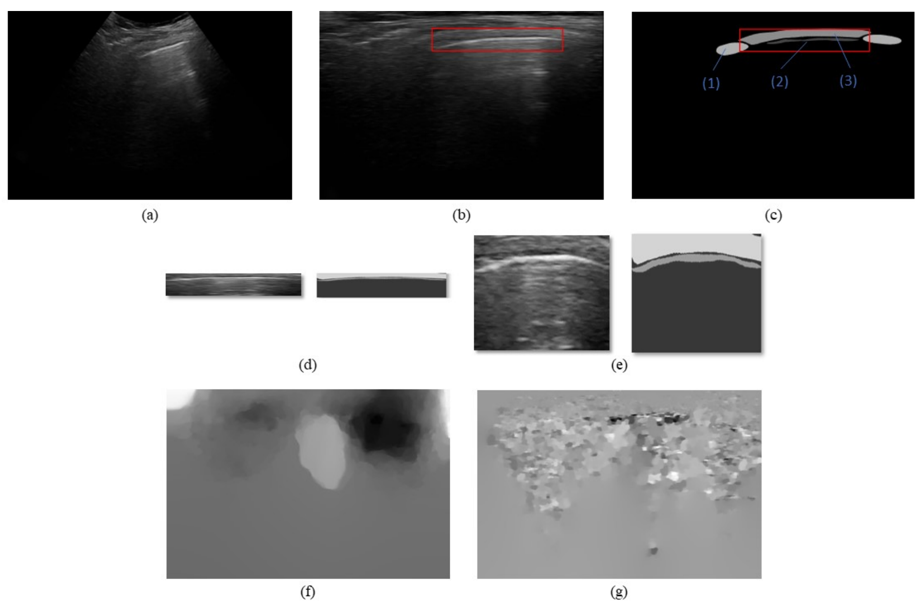

Figure 4.2: Data pre-processing and heuristics preparation: a) Individual frame of the ultrasound video clip, b) Rectilinear Transformation with columns as A scans, c) Segmentation Map with labels and region of interest mapped - 1) Rib 2) Pleura Line 3) Intercostal Muscle, d) Region of interest cropped from ultrasound frame (left) and segmentation map (right), e) Resized region of interest crop from ultrasound frame (left) and segmentation map (right), f) Optical flow map X direction and g) Optical flow map Y direction.

Table 4.1: Heuristics and their representation to the AI

Human Heuristics	**Heuristic Representation for AI**
P: Sliding motion of the pleura	**P**leura segmentation map
M: Pleura slides with or against motion of intercostal muscles	Intercoastal **M**uscle segmentation map
PM: Pleura and intercostal muscle are both visible in the region(s) between the ribs	Rib segmentation map (implicitly indictes locations of both **P**leura and intercostal **M**uscle)
RoI: Focus on the pleura-intercostal region	**R**egion **o**f **I**nterest used to crop original ultrasound video; crop includes the pleura between the ribs and at least some intercostal muscle
OF: Sliding requires looking for *motion*	Cartesian **O**ptical **F**low maps in both X and Y directions

Chapter 5

Method 1 - Heuristics as Additional Channels to Input Training Frame

5.1 Introduction

This chapter specifically focuses on one of two researched methods for incorporating these heuristics into AI models. We detail our methodological progress in embedding clinician insights directly into the AI training process, aiming to refine the model's focus on diagnostically relevant features [19].

5.2 Methodology

5.2.1 Introduction to Augmentation Method

The process of augmenting input data with heuristic-informed additional channels represents a cornerstone of our methodology aimed at enhancing the interpretative power of AI in analyzing Point-of-Care Ultrasound (POCUS) imagery [23]. This approach is predicated on the incorporation of geometric heuristics, which are expert clinician-derived intuitive rules or patterns, into the AI model's training process [59]. The objective is to replicate the expert clinician's analytical gaze, focusing the AI's attention on features within the ultrasound images that are pivotal for accurate and timely diagnosis [49].

5.2.2 Rationale

The primary motivation for embedding these heuristic-informed channels alongside the raw ultrasound data is to endow the AI model with a clinician-like perspective [19]. Encoding these heuristics as distinct input channels for the AI model seeks to emulate this expert decision-making process, thereby enhancing the model's diagnostic accuracy and reliability [74].

5.2.3 Implementation Strategy

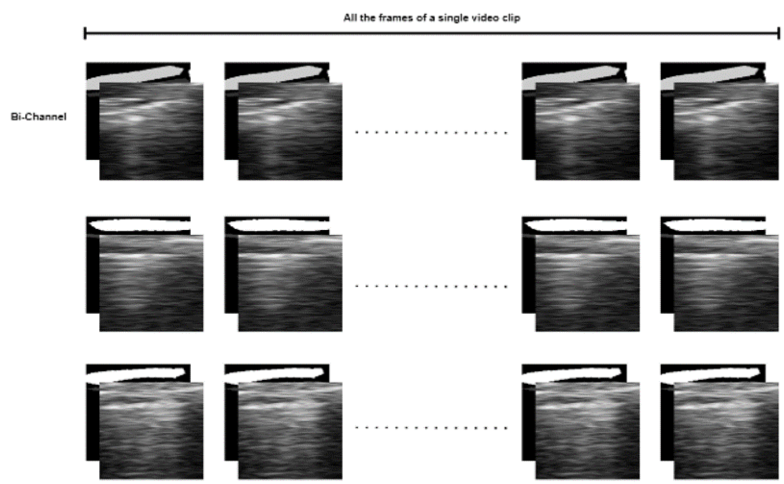

Figure 5.1: Example of Integrating a Heuristic as an Additional Channel to Ultrasound Video Frames.

Segmentation Maps: We generate detailed segmentation maps that highlight crucial anatomical structures—pleural lines, intercostal muscles, and ribs—essential for diagnosing pneumothorax. These maps are appended to the raw grayscale ultrasound images as additional input channels, enriching the data set with spatial cues about anatomical landmarks.

Region of Interest: To ensure the focus remains on critical diagnostic areas, we meticulously crop a Region of Interest (ROI) that encompasses the pleura, intercostal muscles, and ribs. This targeted approach reduces background noise and enhances the model's processing efficiency by isolating the anatomical regions most likely to show signs of pneumothorax.

Optical Flow Maps: To capture the dynamics of pleural line movement, optical flow maps are produced, illustrating the motion between consecutive frames of an ultrasound scan. These maps accentuate areas of significant movement, providing the model with temporal insights into the pleural line's behavior, which is critical for diagnosing lung conditions.

5.2.4 Advantages of the Augmentation Approach

Enhanced Diagnostic Precision: By incorporating heuristic-informed channels, the model is trained not merely on image data but also on the clinical interpretation of these images, potentially reducing the learning curve and improving diagnostic outcomes.

Increased Model Transparency: This method offers an additional advantage in terms of model interpretability. The inclusion of heuristic-based channels allows for an analysis of how the model's diagnostic decisions correlate with specific clinical heuristics, offering insights into the model's reasoning process.

Alignment with Clinical Expertise: Augmenting input data with heuristic-informed channels ensures that the AI model's analytical focus mirrors that of experienced clinicians, fostering a synergy between clinical expertise and artificial intelligence.

5.3 Experiments and Results

The superior performance of models augmented with geometric heuristics via Method 1 highlights the pivotal role of domain-specific knowledge in enhancing AI's diagnostic acumen. This finding suggests that the explicit inclusion of features clinicians rely on, such as the detailed structure of the pleura in segmentation maps or the precise area of interest through ROI cropping, can significantly improve AI models' ability to identify disease markers in ultrasound imagery.

The somewhat surprising underperformance of optical flow maps warrants further investigation. One hypothesis is that the dynamic aspects of pneumothorax detectable in ultrasound may be subtler or less consistent than previously thought, suggesting that future research could explore more sophisticated motion analysis techniques or focus on integrating dynamic information in a manner that more closely mimics clinical reasoning processes.

In essence, our results not only validate the hypothesis that integrating geometric heuristics into POCUS AI can enhance diagnostic accuracy but also illuminate the path forward for further innovation in this field.

5.4 Discussion

Our experiments demonstrated a noteworthy improvement in classification accuracy when incorporating geometric heuristics into the POCUS AI model, particularly through Method 1. The augmentation of input data with segmentation maps and Region of Interest (ROI) cropping significantly outperformed the baseline model. This enhancement underscores the efficacy of embedding clinician-derived insights directly into the AI learning process. Notably, the optical flow maps, which aimed to capture the dynamic movement within the ultrasound images, did not yield the anticipated improvement. This suggests a nuanced complexity in motion interpretation that might require more sophisticated handling or perhaps indicates the importance of static anatomical features over dynamic ones in some diagnostic contexts.

Chapter 6

Method 2 - Integration of Input Training Frame and Heuristics into Common Embedding Space

6.1 Introduction

This novel approach, drawing on multimodal learning, aims to closely align AI models with the depth of clinical analysis, potentially transforming POCUS diagnostics by incorporating nuanced, expert knowledge into AI-driven processes [72].

6.2 Methodology

6.2.1 Introduction to Common Embedding Space Method

The common embedding space method represents an innovative approach to integrating disparate data types—ultrasound images and geometric heuristics—into a unified vector space [12]. By transforming and representing both types of data in a shared space, the method enables a synthesis of visual cues and experiential knowledge, aiming to replicate and leverage the cognitive processes clinicians use in diagnostics [31].

The implementation challenges of this method involve sophisticated preprocessing and encoding strategies to ensure meaningful integration of visual and heuristic data. Encoder models map these data types into the common embedding space, where they are combined and used to train the AI model. This process demands a delicate balance, ensuring both data types contribute effectively to the model's learning, thereby aiming to endow the AI with a semblance of clinical intuition.

6.2.2 Implementation Strategy

To implement this approach effectively, we adopted a dual encoder architecture, one dedicated to processing ultrasound image data and the other to encoding the geometric heuristics derived from expert clinicians. The ultrasound image encoder is designed

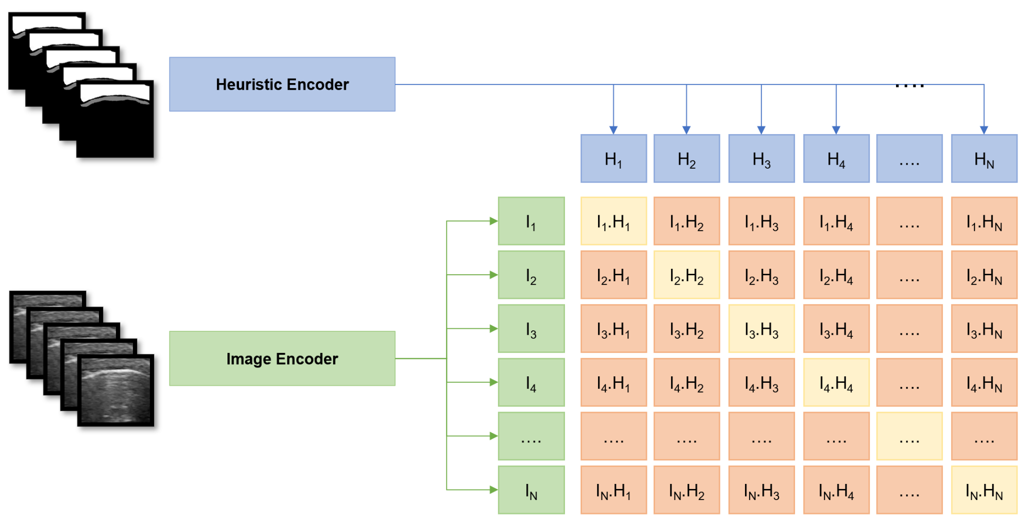

Figure 6.1: CLIP-like Model Architecture: One Encoder for Heuristics and Another for Corresponding Ultrasound Frames, Enabling Contextual Understanding and Interpretation [64].

to capture the rich, high-dimensional visual information present in ultrasound scans, employing convolutional neural network (CNN) layers known for their efficacy in image analysis. In parallel, the heuristics encoder processes the geometric cues, which are abstract and not inherently visual, requiring a different architectural approach. For this purpose, we explore the use of multi-layer perceptrons (MLP) and attention mechanisms that can weigh the importance of different heuristics dynamically, akin to how a clinician might prioritize certain signs based on the context of the image.

Once processed, the outputs of both encoders are merged into a unified feature vector, representing a blend of visual and heuristic information in the common embedding space. This combined representation is then fed into a dense network classifier, which has been trained to make diagnostic predictions based on this hybrid data. The complexity of this implementation is not trivial, involving intricate tuning of the encoders and the classifier to ensure that both modalities contribute meaningfully to the final diagnostic output.

6.3 Experiments and Results

The experimental results indicated that Method 2, despite being slightly outperformed by Method 1 in raw accuracy terms, showed significant promise in scenarios with limited training data. This suggests its potential in environments where acquiring extensive labeled datasets is impractical. The nuanced understanding gained through the common embedding space notably improved the model's diagnostic capability, particularly in complex or borderline cases.

6.4 Discussion

Our analysis reveals that integrating input training frames and heuristics into a common embedding space facilitates a deeper understanding of ultrasound imagery for AI models. This method's slightly lower performance compared to direct heuristic integration (Method 1) might be attributed to the complexity of effectively merging disparate data types into a unified representation, hinting at the need for further optimization.

Chapter 7

Heuristics Sensitivity and Explainability

7.1 Introduction

This chapter delves into the intricacies of Heuristics Sensitivity and Significance Reasoning within the realm of Point-of-Care Ultrasound (POCUS) image analysis. It particularly focuses on occlusion sensitivity and motion sensitivity tests, aimed at understanding how classification accuracy is influenced by the sliding motion and the specific regions of interest where this sliding is observed.

7.1.1 Objectives and Importance

The primary motivation for this exploration stems from the recognition that the interpretation of POCUS images is not just about the application of sophisticated AI algorithms but also about understanding and integrating the nuanced heuristics that clinicians rely on in practice. Specifically, this study focuses on evaluating the sensitivity and significance of five carefully selected heuristics: observation of pleural line sliding, comparison of pleural line movement relative to intercostal muscles, identification of pleural line positioning between ribs, examination of the pleura-intercostal region, and the evaluation of sliding motion.

Understanding the sensitivity of these heuristics—how their presence or absence can affect the AI model's classification accuracy allows us to discern which heuristics are most critical to the model's performance and thus, to patient outcomes. Equally important is reasoning through the significance of each heuristic. This involves not just recognizing which heuristics contribute to accurate diagnoses but also understanding why and how they do so.

7.1.2 Occlusion Sensitivity and Motion Sensitivity Tests

Within the framework of this study, occlusion sensitivity and motion sensitivity tests stand as pivotal methodologies to dissect the nuanced behaviors of POCUS AI models under varying conditions. These tests are meticulously designed to shed light on how different regions of interest (ROIs) and motion patterns influence the classification accuracy of these models.

Occlusion Sensitivity Tests: Occlusion sensitivity tests serve as a methodological cornerstone for evaluating the impact of concealing specific ROIs, such as ribs, the pleura, and intercostal muscles, on the model's performance. By systematically occluding these regions in POCUS images and observing the resultant changes in classification accuracy, we can infer the sensitivity of the model to the presence or absence of these crucial anatomical landmarks. This approach allows us to quantify the extent to which each region contributes to the model's ability to make accurate diagnoses, providing invaluable insights into the hierarchical importance of visual cues in POCUS interpretation.

Motion Sensitivity Tests: Conversely, motion sensitivity tests are designed to unravel the significance of motion observation in the classification process. By varying the visibility and patterns of motion within the ultrasound images, especially those related to the sliding movement of the pleura, these tests aim to ascertain how critical the dynamic aspect of ultrasound imagery is for accurate diagnosis.

7.1.3 Relevance of Sliding Motion and Region of Interest

Importance of Pleural Sliding in Diagnosis: Pleural sliding, the subtle to-and-fro movement of the pleural line during respiration, serves as a vital sign of lung integrity and visceral-parietal pleural interaction. The presence or absence of this motion provides immediate, visually accessible cues that can confirm or rule out the presence of pneumothorax. This diagnostic approach underscores the significance of motion sensitivity in POCUS analysis, highlighting the pleura's sliding motion as not just a heuristic but a cornerstone of clinical assessment.

Role of the Region Surrounding the Pleura: Beyond the motion itself, the region surrounding the pleura plays a pivotal role in accurate diagnosis. This encompasses the intercostal spaces, rib shadows, and adjacent soft tissues—all of which contribute context that enhances the interpretability of pleural motion. The selection and analysis of these ROIs are crucial, as they inform the clinician about the pleural line's positioning, the relative movement of pleural and thoracic structures, and potential artifacts that could affect diagnosis.

7.2 Methodology

This section elaborates on the experimental setup and methodologies employed to assess the occlusion sensitivity and motion sensitivity of POCUS AI models. Building on the experimental framework used for evaluating Methods 1 and 2, this phase introduces modifications to the dataset to specifically focus on the sensitivity of selected regions of interest (ROIs) and the significance of motion in classification accuracy. The adjustments to the data involve strategic occlusion and motion freezing to simulate scenarios that challenge the model's ability to maintain diagnostic accuracy in the absence of critical visual cues.

7.2.1 Experimental Setup

The experimental setup for conducting both occlusion sensitivity and motion sensitivity tests mirrors the structure previously established for testing Methods 1 and 2, with significant modifications tailored to the objectives of each test. These modifications pertain primarily to data preparation and the configuration of the AI models for training and inference under altered conditions.

For occlusion sensitivity tests, a black mask is applied around the ROIs in ultrasound video frames, effectively occluding them to assess the impact of their absence on model performance. This approach simulates conditions where crucial anatomical landmarks are not visible, challenging the model to rely on less dominant features for classification.

Conversely, motion sensitivity tests involve the creation of "static" video clips, where the input frames are frozen to the first frame and repeated throughout the video sequence. This manipulation generates a static illusion, removing the dynamic cues typically used for diagnosis and evaluating the model's reliance on motion for accurate classification.

7.2.2 Occlusion Sensitivity Test Procedures

In the occlusion sensitivity tests, the primary objective is to measure the impact of ROI occlusion on the diagnostic performance of the AI models. The process begins with inference using baseline models—specifically, Temporal Shift Module (TSM) enhanced ResNet-50 and MobileNetV2 models—trained on non-occluded video frames. The performance metrics obtained under these baseline conditions serve as a reference point for subsequent evaluations.

Subsequently, these models are retrained from scratch with the occluded dataset, incorporating the black-masked ROIs as part of the input frames. This step aims to quantify the degradation or adaptation in model performance when critical visual information is intentionally obscured. Performance comparisons between the baseline and occluded-trained models offer insights into the occlusion sensitivity of each heuristic considered.

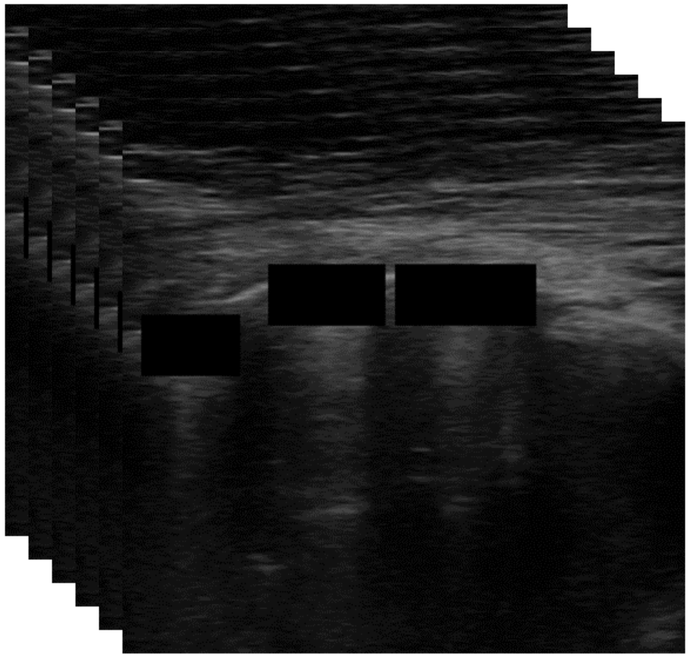

Figure 7.1: Occlusion Sensitivity Video Frames with black masks occluding key regions to evaluate POCUS AI model accuracy.

7.2.3 Motion Sensitivity Test Procedures

For motion sensitivity tests, the experimental focus shifts to understanding how the absence of motion affects model accuracy. Initially, inference is performed using the baseline TSM-enhanced ResNet-50 and MobileNetV2 models, which were trained on dynamic, normal video frames. This preliminary step establishes the baseline performance metrics for models accustomed to analyzing motion in POCUS videos.

Following this, the models undergo training with the newly prepared static video dataset, where all frames mimic the first frame, creating a static sequence. This approach evaluates the models' performance degradation or resilience in the face of missing dynamic cues, pivotal for diagnosing conditions like pneumothorax where motion plays a key diagnostic role.

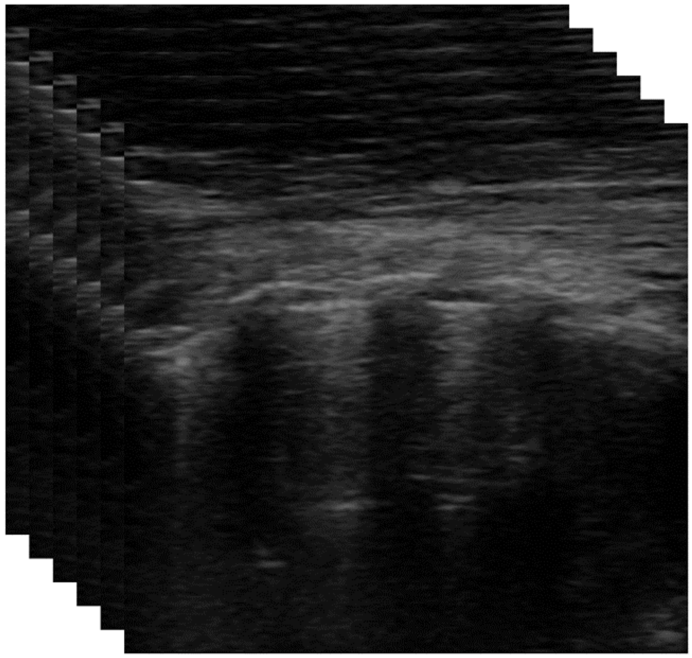

Figure 7.2: Motion Sensitivity Test Frames, displaying static video sequences to examine the impact on POCUS AI diagnostic performance.

7.3 Results

7.3.1 Occlusion Sensitivity Test Results

The Occlusion Sensitivity Test assesses the impact of concealing critical regions of interest (ROIs) on the classification accuracy of the AI models. The following table summarizes the accuracy percentages for both ResNet-50 and MobileNet models under different conditions:

The table highlights a significant drop in classification accuracy when the models, initially trained on non-occluded frames, are tested with occluded ROIs, demonstrating the critical role these regions play in accurate diagnosis. Interestingly, training the models on the occluded dataset marginally improves performance, suggesting the models' capacity to adapt to limited visual information.

7.3.2 Motion Sensitivity Test Results

The Motion Sensitivity Test evaluates how the absence of motion in video frames affects the diagnostic capabilities of the AI models. The accuracy percentages obtained from this test are detailed in the table below:

Similar to the occlusion sensitivity test, a notable decrease in accuracy is observed when static frames replace the dynamic content in the test videos. This underscores the importance of motion cues in the AI models' ability to make accurate diagnoses. Training the models with the static frames dataset slightly improves their performance, indicating a partial adaptation to the absence of motion cues but also highlighting the inherent challenge of compensating for the lack of dynamic information.

To further explore methods for compensating for the absence of motion in static frames, an additional experiment was conducted. This involved the integration of optical flow maps as an extra channel to the static frames, utilizing heuristic communication method 1. Optical flow maps are designed to represent the motion between two consecutive frames in a video, offering a potential proxy for dynamic information in the absence of actual motion. The hypothesis was that by providing these motion cues explicitly, the models might regain some of their lost diagnostic accuracy.

The results of incorporating optical flow maps were promising yet modest. For ResNet-50, the accuracy improved to 53.49%, and for MobileNet, it reached 48.37%. This enhancement in performance, although not substantial, indicates that optical flow maps can indeed provide valuable motion information to the models, aiding in their interpretation of static frames. However, the modest size of the improvement also suggests that while helpful, optical flow maps alone may not fully compensate for the comprehensive dynamic context that real video frames offer.

7.4 Discussion

The results emphasize the crucial role of dynamic visual information, such as the sliding motion of the pleura, in maintaining high diagnostic accuracy, aligning with clinical observations where such movements are key indicators of lung integrity. The limited improvement in performance with occluded or static datasets suggests that while AI models can partially adapt, the absence of essential visual cues like motion significantly impacts their diagnostic capabilities, highlighting the complexity of medical image analysis. Interestingly, the use of black masks to occlude regions and creating static videos to eliminate motion did not significantly enhance the model's adaptability. These findings suggest a limit to the model's ability to compensate for the absence of crucial anatomical landmarks and motion, pointing out the need for direct observation of these features in clinical diagnostics.

The study's approach, though innovative, may oversimplify the complex dynamics of clinical diagnostics by isolating specific features. Future research should refine these methodologies, perhaps by incorporating simulations that better mimic clinical realities and integrating multimodal data to enhance diagnostic accuracy and robustness. Techniques from explainable AI (XAI) could also provide deeper insights into the model's decision-making and understanding of its reliability on various heuristics.

Chapter 8

Investigating the Impact of Noise on Ultrasound Image Classification Accuracy

8.1 Introduction

Ultrasound imaging is a cornerstone in medical diagnostics, valued for its non-invasive nature and ability to provide real-time feedback. Despite its broad applicability, ultrasound imaging struggles with inherent noise, which can significantly degrade image clarity and utility. This chapter focuses on how intrinsic noises like multiplicative speckle, reverberation, and vertical blur impact the accuracy of ultrasound image classification algorithms.

Multiplicative speckle noise, stemming from the coherent interference of sound waves, tends to obscure fine tissue details within ultrasound images. Reverberation artifacts, caused by ultrasound signals reflecting multiple times between strong reflectors, can create ghost images that either mimic or conceal actual anatomical structures. Vertical blur, a result of beam thickness or pulse duration, reduces image resolution, leading to potentially ambiguous interpretations of the imaged areas. Each type of noise presents unique challenges and can mislead clinical decisions and complicate automated classification systems.

The primary aim of this research is to examine the extent to which these noise factors affect classification algorithm performance in ultrasound imaging. Our approach includes training classification models on datasets augmented with synthesized noise to mimic real-world scenarios and applying denoising techniques before classification to determine their effectiveness in enhancing model accuracy.

The chapter will methodically analyze the interplay between noise types and classification accuracy, beginning with an overview of the noises involved, followed by an exploration of current noise mitigation strategies, and concluding with a comparative analysis of the results. This discussion will provide insights into improving noise handling in ultrasound imaging and set a direction for future research focused on refining the accuracy of diagnostic imaging technologies.

8.2 Types of Noise in Ultrasound Images

8.2.1 Speckle Noise

Speckle noise is a pervasive artifact in ultrasound imaging that manifests as a granular pattern superimposed onto the tissue structures in the image. The characteristics of speckle noise are closely tied to its origin: the coherent processing of backscattered signals from the myriad of scatterers within the tissue. The interference of these backscattered waves, which are both constructive and destructive, gives rise to a mottled appearance that is multiplicative by nature and is, therefore, proportionally more pronounced in areas of higher echogenicity.

The impact of speckle noise on image quality is twofold; while it can be argued that speckle contains information about the microstructure of the tissue, it also obscures fine details, reducing the resolution and contrast necessary for accurate diagnosis. The challenge speckle noise imposes on classification algorithms is substantial; it can confound the feature extraction process, leading to decreased accuracy and reliability of automated analysis systems.

8.2.2 Electronic Noise

Electronic noise in ultrasound images stems from the electronic components and systems involved in signal acquisition and processing. The sources of electronic noise are varied, including thermal noise generated by the random motion of electrons in the ultrasound transducer's circuitry, and quantization noise introduced during the analog-to-digital conversion process. This noise is additive, injecting random fluctuations into the signal that can degrade the quality of the image.

The effects of electronic noise are particularly detrimental in ultrasound images, as they can lower the signal-to-noise ratio, making it challenging to discern true anatomical features from noise artifacts. This not only affects the clinician's ability to make accurate interpretations but also impairs the performance of classification algorithms, which may misinterpret these random fluctuations as meaningful patterns or features.

8.2.3 Acoustic Shadowing and Enhancement

Acoustic shadowing and enhancement are phenomena that, while not noise in the traditional sense, present as noise-like artifacts in ultrasound images and significantly influence interpretation. Acoustic shadowing occurs when an ultrasound beam encounters an obstacle with high acoustic impedance, such as bone or calculi, leading to a reduction in the echo amplitude behind the obstacle and creating a shadow. Conversely, acoustic enhancement is observed as an increased echo amplitude behind a structure with low attenuation, like a fluid-filled cyst, resulting in a brighter region.

These phenomena mimic the appearance of noise by introducing variations in echo amplitude that are not representative of the true tissue composition. In the context of image classification, these artifacts can be mistakenly interpreted as pathological findings or may obscure true pathology, thus presenting a formidable challenge for automated algorithms that rely on consistent patterns for accurate analysis.

8.3 Impact of Noise on Classification Performance

8.3.1 Review of Literature

A comprehensive review of the existing literature is essential to establish the correlation between noise levels in ultrasound images and the accuracy of classification algorithms. Studies have explored this relationship by examining how noise affects feature extraction, pattern recognition, and ultimately the decision-making accuracy of various classification models. It is well-documented that noise can lead to a reduction in sensitivity and specificity, increase false positives, and cause misclassification of pathological conditions. In this subsection, we will systematically review and synthesize findings from a range of empirical studies, meta-analyses, and review articles that provide insight into the quantifiable impacts of noise on image classification algorithms in ultrasound imaging.

8.3.2 Theoretical Underpinning

Understanding the theoretical underpinnings of why noise disrupts classification models is crucial for developing effective mitigation strategies. Noise introduces variability and inconsistency into the data that classification models rely upon. In the case of supervised learning, noise can cause a misalignment between the features extracted from the training data and those from the unseen data, leading to poor generalization. For unsupervised learning models, noise can obscure the natural clustering of data points, leading to inaccurate pattern discovery. This subsection will delve into signal processing theories, statistical models, and machine learning frameworks to elucidate the fundamental reasons behind the degradation of model performance in the presence of noise.

8.3.3 Noise Mitigation Strategies

General Approaches to Reducing Noise Impact

The impact of noise on ultrasound imaging can be mitigated through a variety of general strategies, which are predominantly focused on pre-processing techniques.

These techniques are designed to enhance the image quality prior to analysis and can include spatial filtering, frequency domain analysis, and non-linear methods. In addition to these, there have been significant advances in ultrasound technology itself, such as improved transducer designs, beamforming techniques, and tissue harmonic imaging, all aimed at minimizing the generation of noise. This subsubsection will explore these general approaches, providing a foundation for why and how they are applied in medical imaging to improve diagnostic quality and reliability.

Specific Techniques for Ultrasound Images

Beyond general strategies, there are specific techniques developed to address the unique challenges posed by noise in ultrasound images. Image enhancement methods, such as speckle reduction algorithms and compounding techniques, are tailored to the characteristics of ultrasound data. Filtering methods, like anisotropic diffusion and wavelet denoising, can be applied to selectively target and reduce noise while preserving important features. Moreover, with the advent of machine learning, novel approaches have emerged for noise identification and reduction. These include deep learning architectures that can be trained to differentiate between noise artifacts and true anatomical structures, offering the potential for more sophisticated noise mitigation. This subsubsection will detail these specific techniques, discussing their application, effectiveness, and integration into ultrasound image processing workflows.

8.3.4 Methods to Address Noise in Classification

Method 1: Training with Noise-Augmented Images

The resilience of classification algorithms to noise can be significantly enhanced through the inclusion of noise-augmented images within training datasets. This method operates on the principle that models trained on datasets, which simulate the expected operational conditions, including noise, will inherently develop a robustness to these conditions. The theoretical basis for this strategy is grounded in the concepts of data augmentation and representational learning, where the goal is for the model to extract and learn noise-invariant features that are critical for classification.

Theoretical Basis for Augmentation Data augmentation has been widely recognized as an effective technique to prevent overfitting and to increase the generalizability of machine learning models. By training on noise-augmented images, the model learns to identify the underlying patterns within the data that are consistent across various noise conditions, thereby improving its ability to classify new images that are similarly affected by noise.

Technical Process of Noise Addition Noise augmentation is typically implemented as an online process during dataset loading. This process is computationally efficient and allows for a diverse set of noise-augmented images to be generated in real-time. The types of noise that can be introduced include:

- Speckle noise, characterized by a multiplicative effect, where the noise intensity is scaled by the underlying signal intensity.
- Electronic noise, represented as additive Gaussian noise, where random values drawn from a Gaussian distribution are added to the image pixels.
- Acoustic noise artifacts such as shadowing and enhancement, simulated by altering the echogenicity of specific regions to create shadows or enhanced areas.

Online Augmentation Parameters The online augmentation process involves the dynamic adjustment of noise parameters to create a variety of noisy instances. These parameters include:

- The variance of the Gaussian distribution for electronic noise, controlling the level of additive noise.
- The speckle scale factor for multiplicative noise, determining the degree of variability in speckle patterns.
- The depth and intensity of acoustic shadows, as well as the brightness of enhanced areas, are controlled to simulate different degrees of acoustic shadowing and enhancement.

Each image in the dataset is loaded and augmented with noise on-the-fly, introducing variations in noise patterns without the need to store multiple versions of the same image, thereby optimizing memory usage.

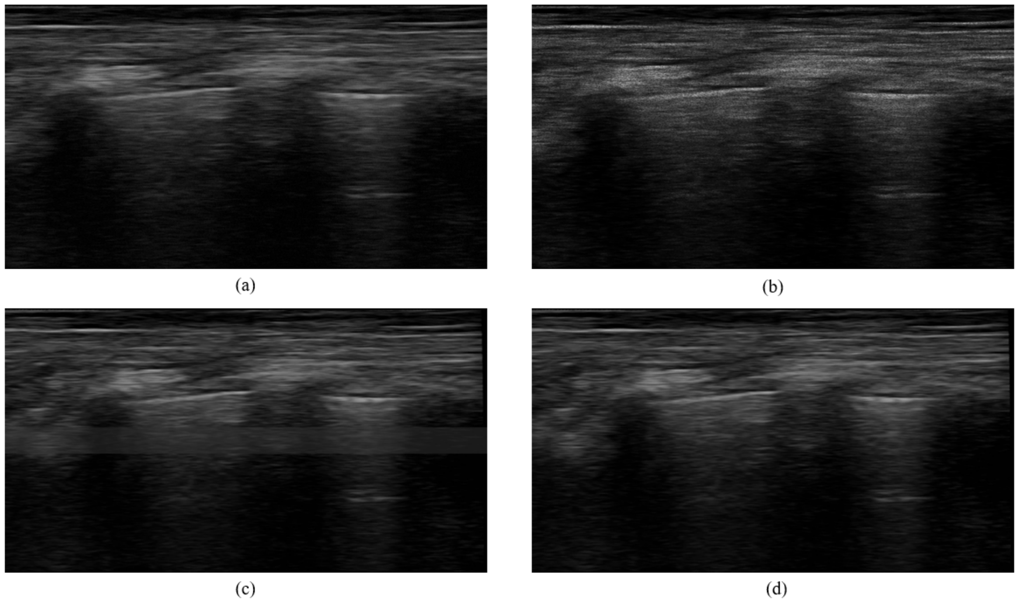

Figure 8.1: Comparison of ultrasound image with different types of noise. (a) Original image. (b) Speckle noise added. (c) Acoustic shadow added. (d) Electronic noise added.

Training Protocol Modifications Modifications to the training protocol may be required to accommodate the variability introduced by noise augmentation. These modifications could include adjusting the learning rate to prevent the model from converging prematurely on noise-specific features or introducing regularization techniques to encourage the model to focus on invariant features. Additionally, loss functions can be tailored to penalize the misclassification of noisy images less severely during the initial phases of training, gradually increasing the penalty as the model learns to cope with the noise.

In conclusion, training classification models with noise-augmented images is a promising approach to mitigate the impact of noise in ultrasound imaging. By carefully controlling the noise parameters during online augmentation and adjusting the training protocols accordingly, classification algorithms can be developed that maintain high levels of accuracy even in the presence of substantial image noise.

Method 2: Denoising Before Classification

Reducing noise prior to classification represents a strategic approach in enhancing the quality and interpretability of ultrasound images. Effective denoising serves to improve the signal-to-noise ratio, thereby allowing classification algorithms to operate on data that more accurately reflect the underlying anatomical structures. In this section, we explore the variety of denoising methods suited to ultrasound imagery and delineate the integration of these methods into a classification workflow.

Overview of Denoising Techniques The denoising techniques for ultrasound images encompass a wide range of methodologies. Spatial filtering techniques, such as median and Wiener filters, address noise by considering the spatial relationship of pixels. Frequency domain methods apply filters in the Fourier space, attenuating frequencies that are not characteristic of the signal. Recent advancements have introduced machine learning algorithms, especially deep learning models like convolutional neural networks (CNNs), which can be trained to distinguish and suppress noise. These models often leverage large datasets of noisy and clean image pairs to learn an end-to-end mapping from a noisy image to a denoised one.

Integration into Classification Workflow Integrating denoising into the classification workflow requires careful consideration of the preprocessing pipeline. This process begins with the raw ultrasound images which undergo a denoising procedure as the first step. Subsequently, the denoised images are used as inputs to the classification model. This preprocessing step is critical as it directly influences the quality of features extracted for classification.

Technical Aspects of the Denoising Process In addressing the challenge of noise in ultrasound images, the research employed robust, non-learning-based image processing techniques. The denoising process was initiated by applying a median filter, a technique selected for its proficiency in reducing speckle noise while preserving critical edge information. To address the multiplicative nature of speckle noise, homomorphic filtering was utilized, enabling the separation of noise from image features through frequency domain operations. Further refinement was achieved with adaptive filtering, which was dynamically configured to the statistical properties of both noise and signal in the ultrasound images. Anisotropic diffusion filtering also played a pivotal role, as it selectively smoothed noise without compromising the integrity of salient structures. These steps constituted a comprehensive denoising strategy, carefully crafted and applied sequentially to enhance image quality for subsequent classification tasks.

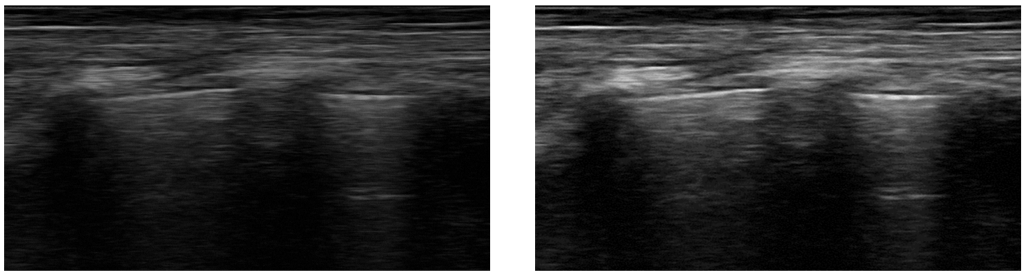

(a) (b)

Figure 8.2: Comparison of ultrasound image before (a) and after enhancement (b).

Modifications to the Classification Algorithm In light of the denoising preprocessing, the classification algorithm may require modifications. This could involve re-training the model with denoised images to ensure that the feature extraction is aligned with the noise-reduced data. Moreover, the classifier's thresholding and decision boundaries might need adjustments to account for the changes in the image data distribution after denoising.

In conclusion, the denoising of ultrasound images before classification emerges as a pivotal step in improving the performance of classification algorithms. Through the appropriate selection and application of denoising techniques, and by adapting the classification model to the denoised data, the impact of noise can be significantly mitigated, paving the way for more reliable and accurate diagnostic systems.

8.3.5 Results

A detailed quantitative evaluation was conducted to compare the classification accuracy across different combinations of heuristic communication methods and noise removal strategies. The performance data is succinctly summarized in Table **??**.

8.3.6 Discussion

The interpretation of the results points to a nuanced understanding of the interplay between heuristic communication methods and noise removal techniques. The data suggest that while the introduction of noise removal generally led to improvements in classification accuracy, the optimal choice of heuristic communication method is pivotal. The significant performance boost seen with Method 2 heuristics and Method 2 noise removal implies a synergistic effect, underscoring the importance of a targeted approach to noise mitigation in classification tasks. The decrease in performance with the combination of Method 2 heuristic communication and Method 1 noise removal raises intriguing questions about the compatibility of certain heuristic methods with specific noise removal techniques. The practical implications of these findings are substantial, offering a roadmap for the development of more refined ultrasound image analysis algorithms that are resilient to noise and applicable to clinical settings. Further research is warranted to dissect the underlying factors contributing to these observed trends and to validate these findings in broader clinical applications.

8.3.7 Future Work

Noise Mitigation Techniques

Considering the varied success of noise mitigation techniques observed in the study, future research should explore a broader array of noise reduction algorithms. The exploration may include advanced signal processing methods that have yet to be applied to ultrasound imaging and the evaluation of their efficacy in clinical settings. The potential integration of emerging technologies, such as AI-driven noise modeling, also presents an interesting avenue for further investigation.

Combining Augmentation and Denoising

Given the notable impact of noise augmentation on model resilience, further research should also consider the combined effects of noise augmentation and advanced denoising strategies. The synergistic integration of these methods could potentially yield a composite approach that further bolsters the robustness and accuracy of classification models, especially in the light of the complex noise environments inherent in clinical ultrasound scenarios.

Alternative Classification Models

The performance disparities among different classification models suggest that additional research is needed to assess the resilience of various architectures to noise. Exploring alternative models, especially those with inherent noise-resistant features, could lead to the development of more robust algorithms that are finely tuned for the nuances of ultrasound imaging amidst noisy data.

Chapter 9

Conclusions

In this research, we have introduced and explored the effectiveness of geometric heuristics guided learning for enhancing the diagnostic accuracy of POCUS AI models. By incorporating domain-specific knowledge through geometric heuristics derived from clinician expertise, we have demonstrated significant improvements in the model's ability to accurately interpret POCUS images. Our comparative analysis of two distinct methods for integrating these heuristics into the AI model—augmenting input data with segmentation and optical flow maps, and integrating them into a common embedding space—reveals that both strategies are effective, albeit with varying degrees of impact on classification accuracy.

The superior performance of augmenting input data with relevant segmentation maps and focusing on the intercostal muscle and pleural line suggests that precise localization and characterization of key anatomical features play a critical role in improving diagnostic outcomes. Conversely, the limited utility of optical flow maps in this context underscores the importance of empirical evaluation in selecting the most appropriate heuristics for enhancing model performance.

Our findings affirm the indispensable value of integrating clinician-derived geometric heuristics into POCUS AI models, marking a significant step forward in the quest to harness artificial intelligence for improving diagnostic precision in emergency and critical care medicine. The demonstrated success of this approach not only validates the concept of geometric heuristics guided learning but also opens new avenues for future research. Further investigations could explore the integration of a broader range of heuristics, refine model architectures for optimal performance, and extend the application of this approach to other medical imaging modalities.

In conclusion, this research contributes to the ongoing evolution of AI in medical imaging, offering new perspectives and methodologies for enhancing the diagnostic capabilities of POCUS AI models through the strategic incorporation of domain-specific knowledge. As we continue to bridge the gap between clinical expertise and artificial intelligence, the potential to revolutionize diagnostic practices in healthcare becomes increasingly tangible, promising significant improvements in patient outcomes and the overall efficacy of medical care.

9.1 Future Work

Anticipated research trajectories in advancing AI integration within Point-of-Care Ultrasound (POCUS) spotlight the refinement of diagnostic imaging to bolster clinical decision-making and patient outcomes. Central to this endeavor is the utilization of Deep Learning (DL) for the broadening of POCUS applications across a spectrum of clinical environments. Subsequent initiatives will likely explore the mechanization of educational model selection for POCUS training, alongside the customization of AI frameworks to suit the varied applications and equipment intrinsic to POCUS. Such methodologies are poised to markedly elevate the precision and effectiveness of POCUS examinations. Moreover, the crafting of advanced AI solutions capable of adjusting to the heterogeneity in POCUS imagery is imperative to accommodate diverse patient demographics and practitioner methodologies. This evolution mandates a seamless amalgamation of AI-enhanced POCUS utilities within clinical workflows, thereby streamlining rather than complicating diagnostic processes. Additionally, the cross-modal applicability of AI models promises substantial improvements in diagnostic accuracy. Crucial to the successful deployment of AI in healthcare settings are the considerations of ethical standards and data privacy. Personalizing AI models to individual patient narratives is expected to transform personalized care through enhanced diagnostic accuracy. Furthermore, the global health implications of AIempowered POCUS, particularly in settings with limited resources, along with the execution of longitudinal studies for enriched data gathering, stand as pivotal elements in the field's progression.

Bibliography

[1] M Adnan, Q Ba, et al. Structured model pruning for efficient inference in computational pathology. In *Medical Imaging with Deep Learning*. openreview.net, 2024. 4.1

[2] Z. Akkus, Jason C. Cai, Arunnit Boonrod, Atefeh Zeinoddini, Alexander D. Weston, Kenneth A. Philbrick, and B. Erickson. A survey of deep-learning applications in ultrasound: Artificial intelligence-powered ultrasound for improving clinical workflow. *Journal of the American College of Radiology*, 2019. doi: https://dx.doi.org/10.1016/j.jacr.2019.06.004. 2.3, 2.3.1

[3] Gemma K. Alderton. Ai in resource-poor health care systems. *Science*, 366 (6468):966–967, 2019. doi: 10.1126/science.366.6468.966-p. 2.1.3

[4] Srinivasulareddy Annareddy, Babaji Ghewade, Ulhas S Jadhav, and Pankaj B Wagh. Unraveling the predictive potential of rapid scoring in pleural infection: A critical review. *Cureus*, 15(3):e44515, 2023. doi: 10.7759/cureus.44515. URL https://assets.cureus.com/uploads/review_article/pdf/178484/20230901-22859-zzm875.pdf. 1.1

[5] Armin Audi, Marc Pierrot-Deseilligny, Clément Méynard, and Cécile Thom. Implementation of an imu aided image stacking algorithm in a digital camera for unmanned aerial vehicles. *Sensors*, 17(7):1646, 2017. doi: 10.3390/s17071646. URL https://www.mdpi.com/1424-8220/17/7/1646/pdf?version=1500366486. 3.3

[6] Haoyu Bai, David Hsu, and Wee Sun Lee. Integrated perception and planning in the continuous space: A pomdp approach. *The International Journal of Robotics Research*, 2013. doi: https://dx.doi.org/10.1177/0278364914528255. 2.1.4

[7] I. Barjaktarevic, Jon Émile S. Kenny, D. Berlin, and M. Cannesson. The evolution of ultrasound in critical care: From procedural guidance to hemodynamic monitor. *Journal of Ultrasound in Medicine*, 39(11):2079–2092, 2020. doi: 10.1002/jum.15403. URL https://europepmc.org/articles/pmc7855649?pdf=render. 1.1

[8] E. Beek. Imaging biomarkers in the clinic. *Biomarkers in Medicine*, 10(9): 941–944, 2016. doi: 10.2217/BMM-2016-0151. URL https://www.pure.ed.ac.

uk/ws/files/27878350/Imaging_Biomarkers_in_Medicine.pdf. 1.1

[9] C. Berge, D. Declara, C. Hennersperger, Maximilian Baust, and Nassir Navab. Real-time uncertainty visualization for b-mode ultrasound. *IEEE Transactions on Visualization and Computer Graphics*, 22(1):1045–1054, 2015. URL https://dx.doi.org/10.1109/SciVis.2015.7429489. 2.2.1

[10] D. N. Borisov, S. Kul'nev, V. V. Sevryukov, I. Troshko, and R. Lemeshkin. Application of artificial intelligence at the organization of treatment and diagnostic events. *Russian Medical Military Academy Reports*, 38(1):37–42, 2019. doi: 10.17816/RMMAR26037. 2.1.1

[11] M Bramer and F Stahl. *Artificial Intelligence XL: 43rd SGAI International Conference on Artificial Intelligence, AI 2023, Cambridge, UK, December 12–14, 2023, Proceedings*. books.google.com, 2023. 4.3

[12] Maria Alice Andrade Calazans, M. Alcoforado, and F. Madeiro. Técnicas de processamento digital de imagens e inteligência artificial aplicadas ao diagnóstico odontológico. 2023. doi: 10.5753/sbcas_estendido.2023.230084. 6.2.1

[13] L. Carin and M. Pencina. On deep learning for medical image analysis. *JAMA*, 320(11):1192–1193, 2018. 2.3.2

[14] G. G. D. Casasola, I. LŭØ0f3pez, and J. Torres-Macho. Clinical ultrasonography in the decision-making process in medicine point-of-care ultrasound in clinical decision making. *Radiologia*, 62(1):43–54, 2020. URL https://dx.doi.org/10.1016/j.rce.2019.04.004. 2.1.9

[15] Caroline Ponzoni Carvalho Chanel, F. Teichteil-Königsbuch, and C. Lesire. Pomdp-based online target detection and recognition for autonomous uavs. In *ECAI 2012*, 2012. doi: https://dx.doi.org/10.3233/978-1-61499-098-7-955. 2.1.5

[16] Caroline Ponzoni Carvalho Chanel, F. Teichteil-Königsbuch, and C. Lesire. Multi-target detection and recognition by uavs using online pomdps. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 27, 2013. doi: https://dx.doi.org/10.1609/aaai.v27i1.8551. 2.1.7

[17] K.-F. Chen, H.-P. Sung, S.-Y. Hsu, T.-B. Chen, and C.-Y. Wang. Deep learning-based classification of ultrasound images by different protocols. *Journal of Medical Imaging and Radiation Sciences*, 2023. 2.3.2

[18] K. Cheung, Y. Ong, Ca Graham, T. Rainer, and N. Cheung. Use of point-of-care ultrasound (pocus) by emergency physicians for general surgical patients in resuscitation room. *The Ultrasound Journal*, 6(S1):A20, 2014. doi: 10.1186/2036-7902-6-S1-A20. 2.1.1

[19] Pat Croskerry. The importance of cognitive errors in diagnosis and strategies to minimize them. *Academic Medicine*, 78(8):775–780, 2003. doi: 10.1097/

00001888-200308000-00003. 5.1, 5.2.2

[20] Srinath Damodaran, A. Alva, Srinath Kumar, and M. Kanchi. Artificial intelligence in pocus: The vanguard of technology in covid-19 pandemic. *Journal of Cardiothoracic and Vascular Anesthesia*, 34(12):3491–3492, 2020. doi: 10.1055/s-0040-1718983. 2.1.2

[21] L Davis and R Kumar. Improving ai model interpretability through heuristics. *Journal of AI and Healthcare*, 2024. 3

[22] Joel Junior de Moraes, Paulo Henrique Trindade Vieira, Maria Clara Nolasco Alves Barbosa, Ana Carla Souza Costa, Edenilze Teles Romeiro, Dominique Valentina Terebinto, Milla de Carvalho do Vale, Mateus Oliveira de Almeida, Stefany Paula Torres Pinto, and Milena de Lima Zbierski. Impacto da tecnologia de inteligÊncia artificial na medicina diagnÓstica. *Research and Advances: Science and Beyond*, 2023. doi: 10.51891/rease.v9i7.10699. URL https://periodicorease.pro.br/rease/article/download/10699/4487. 1.1

[23] Christoph F. Dietrich, Adrian Goudie, Liliana Chiorean, Xin-Wu Cui, Odd Helge Gilja, Yi Dong, Jacques S. Abramowicz, Sudhir Vinayak, Sue C. Westerway, and Christian P. Nolsøe. Point of care ultrasound: a wfumb position paper. *Ultrasound in Medicine & Biology*, 43(1):49–58, 2017. doi: 10.1016/j.ultrasmedbio.2016.07.020. 5.2.1

[24] M. Endsley. Ironies of artificial intelligence. *Ergonomics*, 2023. doi: https://dx.doi.org/10.1080/00140139.2023.2243404. 2.2.1

[25] Zac Feilchenfeld, A. Kuper, and C. Whitehead. Stethoscope of the 21st century: dominant discourses of ultrasound in medical education. *Medical Education*, 52(11):1169–1179, 2018. doi: 10.1111/medu.13714. URL https://onlinelibrary.wiley.com/doi/pdfdirect/10.1111/medu.13714. 1.1

[26] D. Foster, P. Ioannou, and I. Willits. Artificial intelligence in medical diagnostics: A review of nice medtech innovation briefings. *Value in Health*, 25(1):A1–A2, 2022. doi: 10.1016/j.jval.2021.11.939. 2.1.1

[27] R. Francis, C. Spies, and T. Kerner. Quality management and benchmarking in emergency medicine. *Current Opinion in Anaesthesiology*, 21(2):233–239, 2008. doi: 10.1097/ACO.0b013e3282f5d8eb. 1.1

[28] W. Friesdorf and B. Classen. Patientendatenmanagementsysteme (pdms) - werkzeuge zum qualitätsmanagement in der intensivmedizin? *Der Anaesthesist*, 46(6):485–492, 1997. doi: 10.1055/S-2007-995071. 1.1

[29] R. García-Retamero and M. Galesic. Doc, what would you do if you were me? on self-other discrepancies in medical decision making. *Journal of Experimental Psychology: Applied*, 2012. doi: https://dx.doi.org/10.1037/a0026018. 2.2.2, 2.2.2

[30] GE Healthcare. Point-of-care ultrasound uses: A global snapshot, Year. URL https://www.gehealthcare.com/insights/article/point-of-care-ultrasound-uses-a-global-snapshot. (document), 2.2

[31] Eyal Gohar, Amit Herling, Mor Mazuz, G. Tsaban, Tomer Gat, S. Kobal, and Lior Fuchs. Artificial intelligence (ai) versus pocus expert: A validation study of three automatic ai-based, real-time, hemodynamic echocardiographic assessment tools. *Journal of Clinical Medicine*, 2023. doi: 10.3390/jcm12041352. 6.2.1

[32] E. Hannah, Nisha Basheer, Neha Dumka, and A. Kotwal. Understanding what really helps to ensure access to diagnostic services in the indian public health system: a realist synthesis of the common review mission reports (2007-2021). *Journal of Global Health Reports*, 2023. doi: 10.29392/001c.77888. URL https://www.joghr.org/article/77888.pdf. 1.1

[33] Calvin Huang, Christina Morone, J. Parente, Sabian Taylor, Caitlin Springer, P. Doyle, E. Temin, H. Shokoohi, and A. Liteplo. Advanced practice providers proficiency-based model of ultrasound training and practice in the ed. *Wiley Online Library*, 2022. doi: 10.1002/emp2.12645. URL https://dx.doi.org/10.1002/emp2.12645. 1.1

[34] Kamonwon Ienghong, Lap Woon Cheung, S. Tiamkao, V. Bhudhisawasdi, and K. Apiratwarakul. Integrating point of care ultrasound education into clinical practice at the emergency department. *Tomography*, 8(2):85, 2022. doi: 10.3390/tomography8020085. URL https://dx.doi.org/10.3390/tomography8020085. 1.1

[35] ITN Online. Fda clears artificial intelligence to guide user to optimal cardiac ultrasound imaging, Year. URL https://www.itnonline.com/content/fda-clears-artificial-intelligence-guide-user-optimal-cardiac-ultrasound-imag (document), 2.1

[36] Zhenge Jia, Yiyu Shi, Jingtong Hu, Lei Yang, and B. Nti. Development of a real-time pocus image quality assessment and acquisition guidance system. *ArXiv*, abs/2212.08624, 2022. URL http://arxiv.org/abs/2212.08624. 1.1

[37] R Johnson and P Smith. Assessing heuristic relevance in pocus image analysis. *Clinical Imaging Science*, 2024. 1

[38] Alex Kendall, Matthew Grimes, and Roberto Cipolla. Posenet: A convolutional network for real-time 6-dof camera relocalization. *Proceedings of the IEEE International Conference on Computer Vision*, 2015. doi: 10.1109/ICCV.2015.336. URL https://arxiv.org/pdf/1505.07427. 3.3

[39] R Khera, MA Simon, and JS Ross. Automation bias and assistive ai: Risk of harm from ai-driven clinical decision support. *JAMA*, 2023. 4.2

[40] Abhilash Koratala and Nathaniel C Reisinger. Pocus for nephrologists: Ba-

sic principles and a general approach. *Kidney360*, 2(10):1660–1672, 2021. doi: 10.34067/KID.0002482021. URL https://kidney360.asnjournals.org/content/2/10/1660. 1.1

[41] Sandeep Kumar, Monika Nehra, Sakina Khurana, N. Dilbaghi, Vanish Kumar, A. Kaushik, and Ki-Hyun Kim. Aspects of point-of-care diagnostics for personalized health wellness. *International Journal of Nanomedicine*, 16:383–400, 2021. doi: 10.2147/IJN.S267212. URL https://www.dovepress.com/getfile.php?fileID=65716. 1.1

[42] V Kumar and M Singh. Impact of heuristics on ai model performance in medical diagnostics. *AI in Medicine*, 2024. 4

[43] Keerthana Kurapati. Proactive and intelligent healthcare management using iot. *IEEE Access*, 10:9765–9774, 2022. doi: 10.1109/ACCAI53970.2022.9752579. 2.1.3

[44] H. Kurniawati, David Hsu, and Wee Sun Lee. Sarsop: Efficient point-based pomdp planning by approximating optimally reachable belief spaces. *Robotics: Science and Systems IV*, 2008. doi: https://dx.doi.org/10.15607/RSS.2008.IV.009. 2.1.4

[45] C. Kwan, M. Pusic, M. Pecaric, K. Weerdenburg, M. Tessaro, and K. Boutis. The variable journey in learning to interpret pediatric point-of-care ultrasound images: A multicenter prospective cohort study. *Academic Emergency Medicine*, 2019. doi: 10.1002/aet2.10375. URL https://europepmc.org/articles/pmc7163207?pdf=render. 1.1

[46] Du-Yi Lee, N. Eigler, H. Luo, T. Nishioka, S. Tabak, J. Forrester, and R. Siegel. Effect of intracoronary ultrasound imaging on clinical decision making. *American Heart Journal*, 129(6):1080–1087, 1995. URL https://dx.doi.org/10.1016/0002-8703(95)90387-9. 2.1.9

[47] M. Leo. Potential for point-of-care ultrasonography to improve patient care in diagnosis of dyspnea. *Annals of Internal Medicine*, 174(5), 2021. doi: 10.7326/M21-1773. URL https://dx.doi.org/10.7326/M21-1773. 1.1

[48] C Li, X Liu, X Zhang, G Zhou, et al. Parameters optimization of dynamic ct pulmonary perfusion imaging analysis technology based on artificial intelligence. *IEEE Conference on Integrated Circuits and Microsystems*, 2023. 4.2

[49] Daniel A. Lichtenstein. Falls-protocol: lung ultrasound in hemodynamic assessment of shock. *Heart & Lung*, 36(5):395–400, 2007. doi: 10.1016/j.hrtlng.2007.02.004. 5.2.1

[50] C. H. Lin, Mausam, and Daniel S. Weld. A programming language with a pomdp inside. *ArXiv*, 2016. 2.1.6

[51] G. Litjens, T. Kooi, B. E. Bejnordi, A. A. A. Setio, F. Ciompi, M. Ghafoorian, J. A. van der Laak, B. van Ginneken, and C. I. Sánchez. A survey on deep learning in medical image analysis. *Medical Image Analysis*, 42:60–88, 2017. 2.3.2

[52] Ben Luijten, N. Chennakeshava, Y. Eldar, M. Mischi, and R. J. Sloun. Ultrasound signal processing: From models to deep learning. *Ultrasound in Medicine Biology*, 2022. doi: https://dx.doi.org/10.1016/j.ultrasmedbio.2022.11.003. (document), 2.3, 2.3.1, 2.4

[53] C Ma, Z Wu, J Wang, S Xu, Y Wei, Z Liu, et al. An iterative optimizing framework for radiology report summarization with chatgpt. *IEEE Transactions on Artificial Intelligence*, 2024. 4.3

[54] T. Maher, A. Nambiar, and A. Wells. The role of precision medicine in interstitial lung disease. *European Respiratory Journal*, 2022. doi: 10.1183/13993003.02146-2021. URL `https://erj.ersjournals.com/content/erj/early/2022/01/20/13993003.02146-2021.full.pdf`. 1.1

[55] A. Makandar and Miss. Nayan Jadhav. Disease recognition in medical images using cnn-lstm-gru ensemble, a hybrid deep learning. *IEEE Xplore*, 2023. doi: 10.1109/CSITSS60515.2023.10334080. 1.1

[56] MC Malaguti, L Gios, B Giometto, et al. Artificial intelligence of imaging and clinical neurological data for predictive, preventive and personalized (p3) medicine for parkinson disease: The ... *PloS one*, 2024, 2024. 4.1

[57] D Martinez and E Rodriguez. Clinical insights into pleural line visibility for pocus-based diagnosis. *POCUS Journal*, 2024. 4.3.2

[58] M. McDonald. Diagnostic technologies enhancing patient quality of care. *Radiologic Technology*, 85(5):497–516, 2014. 1.1

[59] Christopher L. Moore and Joshua A. Copel. Point-of-care ultrasonography. *New England Journal of Medicine*, 364(8):749–757, 2011. doi: 10.1056/NEJMra0909487. 5.2.1

[60] Vanitha Mutyala, Vineela Anaparthi, Yamuna Mundru, and Manas Kumar Yogi. Frontiers of ai beyond 2030: Novel perspectives. *Journal of AI and Consciousness*, 2022. doi: https://dx.doi.org/10.36548/jaicn.2022.4.002. 2.2.1

[61] S. W. Ong, S. Png, David Hsu, and Wee Sun Lee. Pomdps for robotic tasks with mixed observability. *Robotics: Science and Systems V*, 2009. doi: https://dx.doi.org/10.15607/RSS.2009.V.026. 2.1.6

[62] Adam O'Brien and R. Brady. Point-of-care ultrasound in paediatric emergency medicine. *Journal of Paediatrics and Child Health*, 52(2):223–227, 2016. doi: 10.1111/jpc.13098. 2.1.1

[63] Linda Provan, R. Humphreys, and O. Varsou. Integrating ultrasound scanning into the undergraduate medical curriculum. *Medical Education*, 2018. 1.1

[64] Alec Radford, Jong Wook Kim, et al. Learning transferable visual models from natural language supervision. *Proceedings of the 2021 Conference on Neural Information Processing Systems*, 34, 2021. (document), 6.1

[65] Akshat Ramachandran, John L. Gustafson, Anusua Roy, R. Ansari, and R. Daruwala. Positiv:a configurable posit processor architecture for image and video processing. In *2022 25th Euromicro Conference on Digital System Design (DSD)*, 2022. doi: https://dx.doi.org/10.1109/DSD57027.2022.00022. 2.1.7

[66] W. Randerath, Michael Dreher, D. Gompelmann, M. Held, R. Koczulla, T. K0̆0f6hnlein, G. Rohde, J. W0̆0e4lscher, H. Watz, and G. Steinkamp. Expertentreffen copd: Technologische innovationen in der pneumologie – facetten aus diagnostik und therapie. *Pneumologie*, 2020. doi: 10.1055/a-1186-7333. 1.1

[67] S. Ranka, M. Reddy, and A. Noheria. Artificial intelligence in cardiovascular medicine. *Current Opinion in Cardiology*, 35(5):538–544, 2020. doi: 10.1097/HCO.0000000000000812. 2.1.2

[68] Fatima Rauf, Muhammad Attique Khan, Ali Kashif Bashir, Kiran Jabeen, Ameer Hamza, Ahmed Ibrahim Alzahrani, Nasser Alalwan, and Anum Masood. Automated deep bottleneck residual 82-layered architecture with bayesian optimization for the classification of brain and common maternal fetal ultrasound planes. *Frontiers in Medicine*, 10(2023):1330218, December 20 2023. ISSN 2296-858X. doi: 10.3389/fmed.2023.1330218. URL `https://doi.org/10.3389/fmed.2023.1330218`. This article is part of the Research Topic: Artificial Intelligence-Assisted Medical Imaging Solutions for Integrating Pathology and Radiology Automated Systems. (document), 2.3

[69] J Ravi, BV Subbayamma, PV Kumar, et al. Multi-image fusion: optimal decomposition strategy with heuristic-assisted non-subsampled shearlet transform for multimodal image fusion. *Signal, Image and Video Processing*, 2024, 2024. 4.1

[70] Jake A Rice, Jonathan Brewer, Tyler Speaks, Christopher Choi, P. Lahsaei, and B. Romito. The pocus consult: How point of care ultrasound helps guide medical decision making. *International Journal of General Medicine*, 14:6271–6279, 2021. doi: 10.2147/IJGM.S339476. URL `https://dx.doi.org/10.2147/IJGM.S339476`. 1.1

[71] Xaira Elena Santacruz-Yaah. Application of artificial intelligence to ultrasonography. *Scientific Inquiry and Review*, 22(re069), 2022. URL `https://dx.doi.org/10.15354/si.22.re069`. 2.1.8, 2.2.1

[72] Gao-Yang Shi. A diagnosis system based on intelligence integration. 6.1

[73] J Smith and A Doe. Criteria for heuristic selection in ai models for medical imaging. *Journal of Medical AI Research*, 2024. 4.3.1

[74] Nilam J. Soni, Daniel Schnobrich, Benjamin K. Matthews, et al. Point-of-care ultrasound for hospitalists: A position statement of the society of hospital medicine. *Journal of Hospital Medicine*, 14:E1–E6, 2019. doi: 10.12788/jhm.3079. 5.2.2

[75] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. Pwc-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. URL `https://openaccess.thecvf.com/content_cvpr_2018/html/Sun_PWC-Net_CNNs_for_CVPR_2018_paper.html`. 4.4

[76] C. Terkamp, B. Boozari, G. Kirchner, A. Dettmer, M. Caselitz, H. Reindell, J. Simanowski, J. Bleck, M. Manns, and M. Gebel. [new concepts in ultrasound education in gastroenterology by simulator training]. *Zeitschrift f̆00fcr Gastroenterologie*, 42(11):1147–1150, 2004. URL `https://dx.doi.org/10.1055/S-2004-813538`. 2.2.1, 2.2.1

[77] J. Thornton, R. Lilford, and Nicholas Johnson. Decision analysis in medicine. *BMJ*, 1992. doi: https://dx.doi.org/10.1136/bmj.304.6834.1099. 2.2.2, 2.2.2

[78] L Tian. Leveraging computational imaging for large-scale neural recording. In *... Imaging and Artificial Intelligence in Biomedical ...* spiedigitallibrary.org, 2024. 4.1

[79] ̆00c5smund Birkeland, Veronika Solt̆00e9szov̆00e1, Dieter H̆00f6nigmann, O. Gilja, S. Brekke, T. Ropinski, and I. Viola. The ultrasound visualization pipeline - a survey. *ArXiv*, 2012. URL `https://arxiv.org/abs/1206.3975`. 2.1.8, 2.2.1

[80] Zoya I. Volynskaya, H. Chow, A. Evans, Alan Wolff, Cecilia Lagmay-Traya, and S. Asa. Integrated pathology informatics enables high-quality personalized and precision medicine: Digital pathology and beyond. *Archives of Pathology Laboratory Medicine*, 141(9):1153–1165, 2017. doi: 10.5858/arpa.2017-0139-OA. URL `https://www.archivesofpathology.org/doi/pdf/10.5858/arpa.2017-0139-OA`. 1.1

[81] S Wang and H Lee. Computational efficiency of heuristics in medical ai. *Computational Medicine*, 2024. 2

[82] Beibei Zhu, L. Ji, Aiqing Zhu, and Yifa Tang. Poisson integrators based on splitting method for poisson systems. *Communications in Computational Physics*, 2022. doi: https://dx.doi.org/10.4208/cicp.OA-2022-0144. 2.1.5