

AI-Driven Pneumonia Detection with Interactive Decision Support

Dr. Uddalak Mitra¹, Barnali Das¹, Ritu Kumari¹

IJIS College of Engineering, Kalyani, India

Abstract

Diagnosing pneumonia from chest X-rays is getting tougher these days. Radiologists are juggling more cases and fatigue sets in fast. While deep learning speeds things up, most standard AI models are still “black boxes”—they make predictions without giving any clue on how they got there. That’s a problem in medicine, where trust and clarity matter.

So, here’s what we built: an Explainable AI (XAI) diagnostic platform based on the MobileNetV2 architecture. We took advantage of transfer learning and heavy data augmentation, and after 28 training epochs, the model reached a solid 94.2% validation accuracy.

The best part isn’t just the numbers. We pulled in Grad-CAM to make the whole process transparent—so instead of tossing out a diagnosis and calling it a day, the system creates heatmaps that highlight exactly where the model sees signs of fluid or infection in the X-ray. Now, doctors don’t just get an answer—they can see the evidence behind it, making it a lot easier to trust and double-check the results.

To fit this into real hospital workflows, we made sure everything runs through a secure, local web dashboard built with PHP and Python. Uploading an X-ray is effortless, reviewing the AI’s heatmaps is quick, and you can instantly save or print compliant PDF reports. The design is simple on purpose—no extra hurdles, just useful info.

By blending strong computer vision with open, clear explanations, our platform aims to be a reliable helper, letting doctors stay in control without slowing them down.

Keywords: Pneumonia Detection, MobileNetV2, Grad-CAM, Explainable AI (XAI), Human-in-the-Loop, Healthcare Decision Support Systems

I.INTRODUCTION

Pneumonia is a nasty surprise—it hits fast, clogs up those delicate air sacs in your lungs with fluid or pus, and suddenly, you’re struggling just to breathe. All sorts of pathogens—bacteria, viruses, fungi—can sneak into your lungs and set off this chaos[1] When it happens, those tiny alveoli swell up and can’t transfer oxygen to your blood. You end up short of breath, sometimes dangerously so.

For years, the way to spot pneumonia has been pretty hands-on. Doctors rely mostly on chest X-rays[2]. That means a radiologist—who’s only human—has to squint at endless shades of gray, hunting for those cloudy patches that scream “infection.” It’s a tough gig, especially on a hectic hospital shift. You can picture it: stacks of X-rays, every patient a new puzzle. There’s a jumble of overlapping shadows—ribs, bones, tissues—all making it hard to see those subtle signs of pneumonia. It’s exhausting, both mentally and visually. And when fatigue sets in, mistakes happen. Missed diagnoses mean patients don’t get antibiotics soon enough. For the frail or the very young, those delays can be deadly.

Chest X-rays are still the best tool for diagnosing pneumonia worldwide[3]. They’re cheap, safe, and every clinic has the gear. But finding pneumonia early and reliably is way easier said than done. Even experts slip up—they’re only human.

This is where deep learning, and specifically CNNs, start to change the game[4]. These AI models are insanely good at picking up tiny patterns and textural shifts in medical images that a tired doctor’s eyes might skip right past. They can scan hundreds of X-rays without blinking, never losing focus. But there’s a hitch: they work like mysterious “black boxes”[5]. They spit out a diagnosis without any real explanation for how they got there. For doctors making big decisions, blind trust isn’t an option. They need to see evidence and understand what the AI is actually seeing in the X-ray.

We couldn’t ignore that gap. So we built a new diagnostic platform, engineered from the ground up for transparency. We used MobileNetV2[6]—a super-efficient neural network that doesn’t need fancy, expensive hardware to run—and paired it with Grad-CAM[7], a clever tool that makes AI decisions visible. Instead of just saying “yes, pneumonia,” and moving on, our system creates colorful heatmaps right on top of the original X-ray. These maps show exactly where the AI spotted something odd in the lungs, making its reasoning crystal clear. Doctors can actually see what the machine saw, breaking down the black box and making AI feel more trustworthy in the clinic.

Here’s what we brought to automated respiratory diagnostics:

- We tweaked MobileNetV2 so it runs fast on everyday hospital computers—no pricey gear needed—and still stays sharp with high diagnostic accuracy[9].

- We integrated Grad-CAM natively, so every AI decision comes with a visual explanation. Doctors can check, verify, or push back if the machine’s opinion doesn’t match their own.
- We built a secure, user-friendly web dashboard that links a PHP-MySQL front end with a Python-based deep learning engine behind the scenes. It’s easy to use, keeps patient data safe, and can actually be deployed in hospitals without a tech headache.
- And we didn’t stop there. We developed a module that automatically generates printable PDF reports. These put the original X-ray and the annotated heatmap side by side, making it simple for hospitals to keep thorough, standardized medical records, fully compliant with legal requirements.

By connecting advanced AI with transparency and practical workflow tools, we’re helping the medical field move one step closer to smarter, safer, and more efficient pneumonia diagnosis. This isn’t just about technology—it’s about making every diagnosis clearer, faster, and more reliable for doctors and their patients.

II. LITERATURE REVIEW

A. Overview of Computer-Aided Diagnostic (CAD) Systems

Medical diagnostics used to rely on physical films, but now everything’s gone digital. Tools like PACS and early CAD systems replaced those old-school methods. The first CAD platforms mainly used hand-crafted rules—think edge detection and simple features. They worked, but weren’t reliable and often triggered false alarms. These days, clinicians want systems that are web-based, smart, and fast. With deep learning driving the shift, platforms now deliver accurate, real-time results[10].

B. Evolution of Convolutional Neural Networks (CNNs)

CNNs changed how we analyze medical images. Instead of programmers telling computers what to look for, CNNs learn features on their own at multiple levels—no hand-holding needed. Since AlexNet sparked things in 2012, models have gotten deeper and better. ResNet, for example, tackled the “vanishing gradient” issue with skip connections. Today, when it comes to detecting pneumonia, CNNs are the go-to. They spot everything from normal anatomy like ribs to complicated patterns like fluid collections, all without human help.

TABLE I. SUMMARY OF EXISTING LITERATURE ON PNEUMONIA DETECTION

Study (Year)	Core Methodology	Accuracy	Identified Limitations
Kermany et al. (2018)	InceptionV3	92.8%	Requires a lot of computing power not transparent
Stephen et al. (2019)	Custom CNN	93.7%	Easily overfits to data
Chouhan et al. (2020)	Ensemble Model	96.4%	Consumes excessive memory
Proposed System	MobileNetV2 + Grad-CAM	94.2%	Designed for fast web use fully interpretable

C. MobileNetV2 Architecture and Computational Efficiency

Big networks like ResNet50 need expensive GPUs—not great for small clinics or rural hospitals. MobileNetV2 was built to solve this[6]. It uses inverted residuals (a “narrow-wide-narrow” setup) along with Depthwise Separable Convolutions. In plain terms? It breaks a standard convolution into a simpler spatial layer plus a 1×1 pointwise combo. This trick cuts computation by almost 90%, making MobileNetV2 much faster and perfectly suited for web apps that need to run locally and quickly[9].

D. Explainable AI (XAI) and Grad-CAM Methodologies

A lot of doctors don’t trust AI because they can’t see how it makes decisions—the infamous “black box” problem[8]. That’s where Explainable AI (XAI) steps in. Early methods like CAM were too rigid, but Grad-CAM took over as the standard[7]. It uses gradients from the final convolutional layer to highlight exactly where the AI is looking in a medical image. For a doctor, this means they get a heatmap that shows the AI is focusing on actual lung problems (not just random image junk like wires or tubes)[2].

E. User Interface (UI) and Security Compliance

Clunky interfaces frustrate healthcare workers and make them miss important alerts. You need something clean, user-friendly, and clear—color-coding helps a lot when things get busy. Security’s non-negotiable, too. Patient info has to stay private. The system uses a “local-first” setup with XAMPP, so data never leaves the clinic’s network. Everything runs through secure MySQL databases and parameterized queries, keeping X-ray

data locked down and safe from both inside and outside threats.

III. METHODOLOGY

Let me take you behind the scenes of how we turned our pneumonia detection AI from a promising concept into something doctors can actually put to work every day. We didn't want to stop at a smart mathematical model that just sits in a research paper. This was about making something real—a tool that fits naturally into hospital life and genuinely helps, not adds to the chaos.

From the very start, every step of our Software Development Life Cycle (SDLC) revolved around two basic questions: “Does this actually help clinicians?” and “Can this work in the real world, not just in a lab?” Building on that, we dove in—getting our hands dirty with messy, real-world data, picking a neural network architecture that doesn't need a supercomputer to operate, and finally wrapping everything in a web interface simple enough for busy medical staff to use. We kept asking ourselves: “Will a nurse, resident, or radiologist be glad to have this, or will it just gather dust on a server?” We knew our platform had to fit neatly into daily hospital routines—like an old friend, not an intrusive know-it-all.

First, we took raw radiographic data—thousands of chest X-rays[16]—and put it through a rigorous cleaning process. It wasn't enough to just collect pictures; we meticulously labeled them, checked everything twice, and made sure we weren't cheating by letting our training set “leak” into our test set. The AI itself is no slouch: we used transfer learning with proven neural network backbones[1], pushed our training by augmenting images in real time[17], and connected the model to a responsive web dashboard. The result isn't just software; it's a complete ecosystem, one that makes sense in the real world, not just at a conference poster.

A. Data Collection and Preprocessing Pipeline

Anyone who's dealt with medical data knows that it can be a mess—there's noise, artifacts, random errors, and details that no textbook ever mentions. We weren't looking for “ideal” X-rays; we wanted the real deal. So we built a huge dataset containing 5,856 frontal-view chest X-rays, covering everything from crystal-clear pneumonia to the tiniest, most debatable hints of disease[16]. Our dataset wasn't limited to textbook patterns. Some images showed textbook pneumonia, but a lot were subtle, challenging—even seasoned radiologists might argue about them. And the messiness of hospital life was front and center: variable machine exposures, slight patient motion, oxygen tubing in the frame—anything a doctor really sees, we kept in.

We didn't just label the images as “Pneumonia” or “Normal.” We added rich clinical context: patient IDs, biological indicators like fluid status, and technical info about each image. This wasn't just for completeness—it gave us a safety net for data integrity. By maintaining strict train/test splits and checking labels repeatedly, we made sure our final accuracy numbers actually meant something. No “leakage,” no trickery.

X-rays come in all kinds of sizes and shapes, so we had a lot of work to do before the images were AI-ready. We resized every scan to a standardized 150x150 pixels, forced them into a three-channel RGB format, and used Min-Max normalization to scale pixel intensities between 0 and 1. To make sure our AI didn't just memorize what it saw, we threw real-time data augmentation at it—random flips, rotations, shears, and zooms made the system truly robust[17]. This way, even if tomorrow's X-ray is off-center or taken on a different machine, the AI has already “seen” something like it.

Here's the big picture of our data pipeline:

- Context-rich labeling: Patient info, image size, and clinical metadata
- Capturing subtle biological clues: Not just “pneumonia or not,” but the shades in between
- Data integrity: Isolated train/test splits, with double-checked clinical labels
- Consistent image processing: All images become 150×150×3 tensors, pixel-normalized
- Real-time augmentation: On-the-fly rotations, flips, and zooms to boost generalization
- The result: A tough, real-world binary classification challenge—Normal vs. Pneumonia

B. Deep Learning Algorithms and Architecture

Spotting pneumonia on a chest X-ray isn't easy, even for experts. Sometimes the signs jump out, but often they're buried in shadows and subtle patterns. Speed and efficiency mattered for our platform—most hospitals can't afford huge servers or wait minutes for a result. That's why we steered clear of big, computationally heavy models and picked MobileNetV2[6]. It's slim, crazy-fast on everyday computers, and still delivers excellent accuracy.

We leaned on Transfer Learning for a head start[1]. Instead of training a network from scratch, we gave ours a jump by loading pre-trained weights. The core layers already knew the basics—edges, shadows, geometry. That let us focus training on what matters for pneumonia and build our own classification head on top.

Hyperparameter	Configured Value
Optimizer	Adam (Adaptive Moment Estimation)
Learning Rate	0.0001
Batch Size	32
Maximum Epochs	28
Loss Function	Binary Cross-Entropy

Let's break down the architecture:

- MobileNetV2 base: Uses "depthwise separable convolutions" for speed, processing images intelligently without wasting resources. It's sharp enough to spot anything from distinct bone outlines to faint hazy opacities.
- GlobalAveragePooling2D: Flattens the data so the network learns the big picture, not just pixel-level quirks. This helps avoid locking onto irrelevant tiny regions.
- Dense layers with Sigmoid activation: These layers handle more complicated decision boundaries, and that final sigmoid neuron gives a clear, confidence-based yes/no answer.
- Adam Optimizer: Keeps training on track, adjusting neural weights quickly and accurately to reduce error as fast as possible.

TABLE II: LAYER CONFIGURATION OF THE PROPOSED MODEL

Layer (Type)	Output Shape	Param #	Activation
MobileNetV2 (Base)	(None, 7, 7, 1280)	2,257,984	None (Frozen)
GlobalAvgPool2D	(None, 1280)	0	N/A
Dense (Hidden)	(None, 128)	163,968	ReLU
Dense (Output)	(None, 1)	129	Sigmoid

C. Explainable AI & Human-in-the-Loop Integration

Raw accuracy isn't enough in healthcare. Doctors have to trust the tool before they'll use it. That's why we baked transparency right into our system[8]. No doctor wants to rely on a black-box algorithm that just spits out an answer with no explanation.

Everything runs inside a secure, local environment: the AI model, a PHP-based web dashboard, and a MySQL database. Every diagnostic step gets logged and can be traced back at any time.

TABLE III: HYPERPARAMETER TUNING AND OPTIMIZATION CONFIGURATION

Here's how a real-life review session works:

- Step One: The dashboard first shows the unedited X-ray, exactly as it was captured—no tricks, no filtering. Doctors get to form their own first impression.
- Step Two: We generate a Grad-CAM heatmap right next to the image[7]—this overlay lights up the regions of the lungs that the AI found most important. Doctors can see specifically what drew the AI's attention. No ambiguity here; it's obvious if the system is focused on a patch of fluid or a suspicious opacity.
- Step Three: The AI's confidence score (straight from the output neuron) is visible—doctors see both the thumbs-up/thumbs-down and how sure the algorithm is.
- Step Four: When doctors finish their review, they can download a ready-made PDF report. The tool gathers patient details, the original X-ray, the heatmap, and the AI's prediction—everything bundled together, easy to share, or archive.

TABLE IV: MODEL DETAILS AT EPOCH 28

Metric	Performance Value
Peak Training Accuracy	96.48% (Epoch 28)
Validation Accuracy	94.2% (Epoch 28)
Test Accuracy	About 96%
Loss (Binary Cross-Entropy)	0.082
Optimizer	Adam (Learning Rate: 0.0001)

And through all of this, the clinician remains in charge. Our AI acts as an extra set of tireless, well-trained eyes—making suggestions, not commands. The goal was never to replace doctors but to back them up and slot seamlessly into the world they already know, not force them to adapt to yet another complicated piece of technology.

In the end, what we built isn't just smart software—it's a real, transparent, and ethical partner for clinicians. It lets them keep control, trust the process, and—most importantly—spend their time and attention where it really matters: with their patients.

IV.RESULTS

The first phase of our project was all about proving, without a doubt, that our approach worked at a core technical level. We put the optimized MobileNetV2 model through the wringer, pushing it to spot even the faintest traces of pneumonia in a massive set of chest X-rays filled with real-world variation—different patient bodies, different ways the X-rays were taken, and all kinds of noise.

A. Deep Learning Model Performance

To figure out if these numbers actually meant anything, we hit the model with a fresh batch of X-rays it hadn't seen before. Overfitting? We didn't let it take over. We followed a classic split — 80% of the samples went to training and tuning, and the other 20% stayed locked down just for checking results as we went.

For 28 training rounds, we kept checking accuracy and Binary Cross-Entropy loss. The model got sharper with every epoch. Training accuracy jumped up to 96.48%, and validation accuracy stuck close at 94.2%. That gap was tiny, and honestly, that told us the network was actually paying attention to real signs of disease, not just memorizing the data. When we let it loose on totally new cases, it still hit about 96% accuracy [10].

At first though, things weren't as smooth. The model struggled because the data was a bit lopsided — lots more X-rays of infected lungs than healthy ones [1]. So, we pumped up the data augmentation with random rotations, shears, zooms, and flips [17]. This forced the model to deal with all sorts of quirks and helped it zero in on actual abnormalities, not just weird angles or differences in equipment.




Fig. 1. Training and validation accuracy curves of the optimized MobileNetV2 architecture over 28 epochs.

V.DISCUSSION

After locking down the scientific side, we switched gears for the second phase and put ourselves in a clinician's shoes. It's not enough to crunch numbers behind the scenes or spit out a yes/no answer. Doctors need to see why a decision was made—they deserve clarity, not another confusing black box [8]. We aimed for a system that doctors could trust and actually use right there in the clinic.

A. Computational Efficiency and Accessibility

By the end, what we built wasn't just a high-performing algorithm. We designed a tool that hit the sweet spot between accuracy and speed. Using MobileNetV2 meant we didn't need a supercomputer [6]. The model ran its predictions in seconds on everyday machines. That kind of practical, accessible diagnostic ability can be a game-changer—especially for rural hospitals and resource-strapped clinics desperately in need of dependable AI support, without all the heavy infrastructure or expense.

B. Clinical Decision Support: Web Inference and Grad-CAM

We knew the dashboard needed to be simple and direct. For every new X-ray, the system not only labels it (normal or pneumonia) but also highlights the relevant features it picked up on. The magic here is Grad-CAM [7]. It creates a visual map showing exactly which parts of the lung influenced the model's verdict [2]. Instead of leaving doctors guessing, they can see the "thought process" right on the image. With every case, the dashboard shows a prediction label, a confidence score (like 99.8%), and a visual heatmap overlay where green means Normal, and red flags Pneumonia. When a score hovers closer to the decision boundary, the case is probably nuanced and deserves a closer look. This setup puts clinicians in the driver's seat, combining automation with flexible judgment.



Fig. 2. The interactive clinical dashboard interface showing a real-time pneumonia detection analysis with 99.7% confidence.

C. Comparing Clinical Workflow and Reporting

We wanted transparency at every turn so doctors can cross-check the model's output against their own judgment and experience. To keep everything streamlined, predictions, patient info, and model confidence all go straight into a MySQL database on the local XAMPP server. Instead of filling the database with heavy image files, we just store file paths to keep things fast and efficient.

Every time the system finishes an analysis, it automatically creates a PDF report using PHP and the FPDF library. The report shows both the original chest X-ray and the Grad-CAM heatmap side by side, along with a handy legend. Doctors just hit download, drop the file into the patient's record, and move on. From upload to PDF download, the whole process only takes a few seconds.




Fig. 3. Sample of the automated clinical PDF report featuring side-by-side comparative analysis of the raw chest X-ray and the localized Grad-CAM heatmap.

VI. CONCLUSION

Deep learning isn't just tweaking how respiratory diagnostics happen—it's flipping the whole script. Medical professionals aren't stuck squinting at X-rays by themselves anymore. Now, algorithms do the heavy lifting, running calculations and spotting patterns that might slip past even the keenest human eye[4]. In our study, we really put that to the test. We took MobileNetV2—a model already known for its efficiency—and tuned it using transfer learning[1], then

threw in plenty of real-time data tweaks to make it even sharper[17]. The results? Almost 96.5% training accuracy, and the validation accuracy sat close behind at 94.2%. That tight gap tells you everything you need to know: the model's not just parroting back the training data. It actually knows how to handle fresh, real-world cases.

What's even more important is how accessible we've made this. Instead of demanding banks of high-powered GPUs or deep pockets for server time, our system runs just fine on a typical office computer[9]. A regular desktop in a local clinic? Good enough. We wanted to make sure clinics—and not just those in big, well-funded hospitals—could put this tool to work, even with limited technical staff or tight budgets.

But here's the thing: accuracy sitting on a lab's shelf doesn't help patients or doctors in a busy ward. Deep learning models are notorious for being black boxes—sometimes they just spit out results, and nobody can say why. That's not good enough when decisions involve a person's life. To open up that black box, we built an interactive, web-based dashboard with a full suite of explainable AI tools[8]. For each X-ray, Grad-CAM visualizations kick in, literally spotlighting the murky spots on lung images[7]: regions that look extra cloudy, fluid build-ups, strange shadows that might signal something serious. Instead of dumping a probability score and calling it a day, the dashboard overlays clear, color-coded heatmaps right on the scan, along with confidence numbers. It's evidence you can see—not just trust. Suddenly, doctors don't have to guess why the AI said what it did. They get concrete visuals they can interrogate and discuss with colleagues or even with patients' families.

We designed this system to help—not sideline—the experts. It's not here to replace doctors, but to offer real, stress-relieving support when their workload piles up. The dashboard slides right into their routines, without forcing them to learn a dozen new workflows or deal with admin headaches. Radiologists spend less time hunched over screens and more time focusing on tricky cases. Fast triaging becomes possible, and standardized, easy-to-read PDF reports automatically sum up the findings, showing both the untouched X-ray and the highlighted suspicious spots for comparison. Since it's all built on lightweight tech (just PHP and MySQL), it doesn't care if it's in a city hospital or a rural clinic with spotty internet—it fits right in, reliably and smoothly.

And we're just scratching the surface. The way the system's built, it's ready to level up. With a bit more training, it can learn to tell apart not just pneumonia, but what kind—bacterial, viral, fungal, you name it. We know there's a massive need to spot things like COVID-19, tuberculosis, and other tricky lung conditions. Our model's architecture allows for these planned updates.

Plus, there's a clear path to move everything to a cloud-based web app and API, which opens doors for remote clinics or developing areas. When specialist doctors are hours—or even days—away, instant, trustworthy diagnostics from the cloud can genuinely change outcomes for patients.

In the end, this isn't just about math or computer vision. It's about creating tools that work for people: understandable, efficient, and ready to scale wherever they're needed. We believe this is a crucial step toward a future where artificial intelligence empowers, rather than intimidates, both medical staff and the communities they serve. Whether it's in big-city ERs or one-room clinics, this kind of AI doesn't just promise better diagnostics; it makes them possible for everyone.

References

- [1] F. Alshanketi, A. Alharbi, M. Kuruvilla, V. Mahzoon, S. T. Siddiqui, N. Rana, and A. Tahir, "Pneumonia detection from chest X-ray images using deep learning and transfer learning for imbalanced datasets," *Journal of Imaging Informatics in Medicine*, vol. 38, no. 1, 2024.
- [2] J. Colin and N. Surantha, "Interpretable deep learning for pneumonia detection using chest X-ray images," *Information*, vol. 16, no. 1, p. 53, 2025.
- [3] N. Darapaneni, A. Ranjan, D. Bright, D. Trivedi, K. Kumar, V. Kumar, and A. R. Paduri, "Pneumonia detection in chest X-rays using neural networks," *arXiv preprint arXiv:2204.03618*, 2022.
- [4] D. S. Kermany, M. Goldbaum, W. Zhang, and K. Zhang, "Identifying medical diagnoses and treatable diseases by image-based deep learning," *Cell*, vol. 172, no. 5, pp. 1122–1131, 2018.
- [5] I. D. Mienye, "Deep convolutional neural networks in medical image analysis: A review," *Information*, vol. 16, no. 3, p. 195, 2025.
- [6] M. Sandler, A. Howard, M. Zhu, A. Zhmoginov, and L. C. Chen, "MobileNetV2: Inverted residuals and linear bottlenecks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 4510–4520, 2018.
- [7] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, "Grad-CAM: Visual explanations from deep networks via gradient-based localization," in *Proceedings of the IEEE International Conference on Computer Vision (ICCV)*, pp. 618–626, 2017.
- [8] M. T. Zamir, "Explainable AI-driven analysis of radiology reports using text and image data: Experimental study," *JMIR Formative Research*, vol. 9, e77482, 2025.
- [9] M. S. A. Reshan, K. S. Gill, V. Anand, S. Gupta, H. Alshahrani, A. Sulaiman, and A. Shaikh, "Detection of pneumonia from chest X-ray images utilizing MobileNet model," *Healthcare*, vol. 11, no. 11, p. 1561, 2023.
- [10] D. Zhang, F. Ren, Y. Li, L. Na, and Y. Ma, "Pneumonia detection from chest X-ray images based on convolutional neural network," *Electronics*, vol. 10, no. 13, p. 1512, 2021.
- [11] I. Sirazitdinov, M. Kholiavchenko, T. Mustafaev, Y. Yixuan, R. Kuleev, and A. Ilin, "Deep neural network ensemble for pneumonia localization from a large-scale chest x-ray database," *Computers & Electrical Engineering*, vol. 78, pp. 388–399, 2019.
- [12] D. Saranyaraj, "PneuNet: A lightweight convolutional neural network with multiscale feature fusion for automated pneumonia detection from chest X-rays," *Journal of Digital Imaging*, vol. 38, no. 2, 2025.
- [13] D. S. Kermany, M. Goldbaum, W. Zhang, and K. Zhang, "Identifying medical diagnoses and treatable diseases by image-based deep learning," *Cell*, vol. 172, no. 5, pp. 1122–1131, 2018.
- [14] O. Stephen, M. Sain, U. J. Maduh, and D. U. Jeong, "An efficient deep learning approach to pneumonia classification in healthcare," *Journal of Healthcare Engineering*, 2019.
- [15] V. Chouhan et al., "A novel transfer learning based approach for pneumonia detection in chest X-ray images," *Applied Soft Computing*, vol. 98, 106732, 2020.
- [16] X. Wang, Y. Peng, L. Lu, Z. Lu, M. Bagheri, and R. M. Summers, "Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 2097–2106, 2017.
- [17] C. Shorten and T. M. Khoshgoftaar, "A survey on image data augmentation for deep learning," *Journal of Big Data*, vol. 6, no. 1, pp. 1–48, 2019.

Acknowledgments and AI Disclosure

We appreciate Google Gemini and other LLM tools for helping with technical documentation, formatting, and tidying up the manuscript. AI assisted especially in optimizing the Python-to-PHP logic and making sure we met IJSREM's guidelines around explainability and clinical reporting. The core research, model optimization, data work, and testing were handled by the human authors. All AI-generated material has been checked for accuracy and ethical standards.