

# Bone Fracture Detection and Localization using Deep Learning

Venkat koushiik R , Irshad basha S, Deepa H R , Amarnath S S  
RV Institute of Technology and Management, Bengaluru, Kothnur, Karnataka

**Abstract** - Bone fractures are pretty common, and if they aren't caught early, they can mess up recovery. Usually, doctors check X-ray images to spot fractures, but doing it manually takes time and they might miss small or tricky ones. So, to speed things up and improve accuracy, our project uses deep learning for bone fracture detection and localization. We've used the Faster R-CNN model with the Detectron2 framework to automatically detect different types of fractures like wrist, elbow, and forearm. The model is trained on X-ray images with annotations, so it can draw bounding boxes around the fracture area. We've also connected it with a Flask API to give quick, real-time results. This system helps as a second opinion and can reduce errors while saving time. Overall, it shows how deep learning can actually be useful in healthcare diagnostics.

**Keywords:** Bone fracture, Deep learning, Faster R-CNN, Detectron2, X-ray, Flask API

## 1 REVIEW

### 1.1 Introduction

So, we're literally in a time where data is being thrown at us from everywhere. Like seriously, every device, system, or app is generating tons of data non-stop. Whether it's from engineering stuff, social media, e-commerce, healthcare, or even bio-research, there's data pouring in like crazy. Especially with all these digital gadgets, the speed at which digital data is growing is actually insane. Back in 2011, studies said the data volume had increased 9x in just five years, and by 2020, they predicted it would hit 35 trillion gigabytes. That's where the whole term "Big Data" came from—because we just needed a way to describe this huge data explosion.

There've been a bunch of papers and reports trying to explain what "Big Data" really means. Some focus on the challenges and opportunities it brings, others on the background or tech platforms used to deal with it. One big overview by McKinsey Global Institute even broke it down into three main angles—innovation, competition, and productivity. Apart from that, people have also looked into Big Data in more specific contexts. Like,

there's been work on how it applies in IoT (Internet of Things), or how it's being handled in new-gen wireless networks like 5G.

And yeah, a lot of researchers have also tried designing new models and algorithms from a data mining perspective to tackle this mess.

Now coming to our field—machine learning has been a total game-changer over the last decade. It's been helping out in fields like medicine, astronomy, biology, and more, mostly because it's great at picking out useful patterns from huge sets of data. But here's the thing—as Big Data started taking over, traditional ML methods started falling short. Those old-school models expected the entire dataset to fit in memory and work in a very controlled way. But with

this kind of massive data, that just doesn't cut it anymore. So even though Big Data has so much potential to unlock better science and engineering, we clearly need new ways to handle it smartly.

Right now, we're in this crazy era where digital data is being generated literally every second from all over—phones, medical machines, social media, everything. And when it comes to medical stuff, especially with things like X-rays and other scans, the amount of data doctors and hospitals deal with is on another level. That's actually what led me to take up this project on detecting bone fractures using deep learning—it's just a solid example of how we can use all this data to actually help people.

So here's the thing: doctors have tons of fracture cases coming in daily, and going through all the X-rays manually takes time and can lead to errors.

With the way medical data is growing, traditional ways of analyzing it just can't keep up. That's where deep learning comes in. It's basically giving us a smarter, faster, and more accurate way to find and pinpoint fractures in bone X-rays—without depending fully on human eyes.

In the past few years, there's been a lot of buzz around using AI and machine learning in healthcare. Researchers have already applied it in stuff like cancer detection, brain scans, and even predicting diseases. Now, with this project, I'm exploring how we can bring that same power to orthopedics—by training models to look at bone fractures

and identify exactly where the damage is.

But this isn't just about automation. It's about improving accuracy and reducing diagnosis time. Especially in emergency cases, every second matters, and if a system can tell a doctor, "Hey, there's a fracture here," it could seriously help save time and maybe even lives. Plus, with models like Faster R-CNN (which I used with Detectron2 in my project), the goal is not just to say "Yes, there's a fracture," but to actually localize it—meaning show exactly where the fracture is on the image. So yeah, this project is my small step toward using deep learning not just as a cool tech tool but as something that can actually make a difference in how fractures are spotted and treated.

For example, deep learning models like CNNs and RNNs have been game-changers, especially in tasks like image and speech recognition. Transfer learning helps when we don't have a lot of labeled data by reusing knowledge from pre-trained models. Active learning focuses on choosing the most useful data points to label, which saves time and effort. And when datasets are too large to train on a single machine, distributed learning makes it possible to train across multiple systems in parallel. Each of these methods has its strengths, and combining them often leads to even better results.

### *1.2 Quick Look at Machine Learning Techniques*

So before I explain how I used deep learning in this fracture detection system, let's go over what machine learning really is. Basically, ML is all about building systems that can learn from data without being explicitly told what to do every time.

#### *1.2.1 Definition and classification of machine learning*

So before I explain how I used deep learning in this fracture detection system, let's go over what machine learning really is. Basically, ML is all about building systems that can learn from data without being explicitly told what to

do every time. It's used across all kinds of fields—AI, stats, optimization, and even stuff like cognitive science and control systems. In short, it's everywhere. In fact, ML has already shown up in many areas—like YouTube recommendations, Google Translate, self-driving tech, and even spam filters. And now it's pushing into medical domains too, like reading MRIs, analyzing ECG signals, or in my case—spotting bone fractures in X-ray images.

here are mainly three types of learning you'll come across Supervised Learning: This is what I used in my project. It's like training your model with examples—"This is a wrist fracture," "This is normal," etc. You give it labeled data and it learns from it.

Unsupervised Learning: No labels here. The system just tries to find hidden patterns. Could be useful in clustering patient data but not for my use-case.

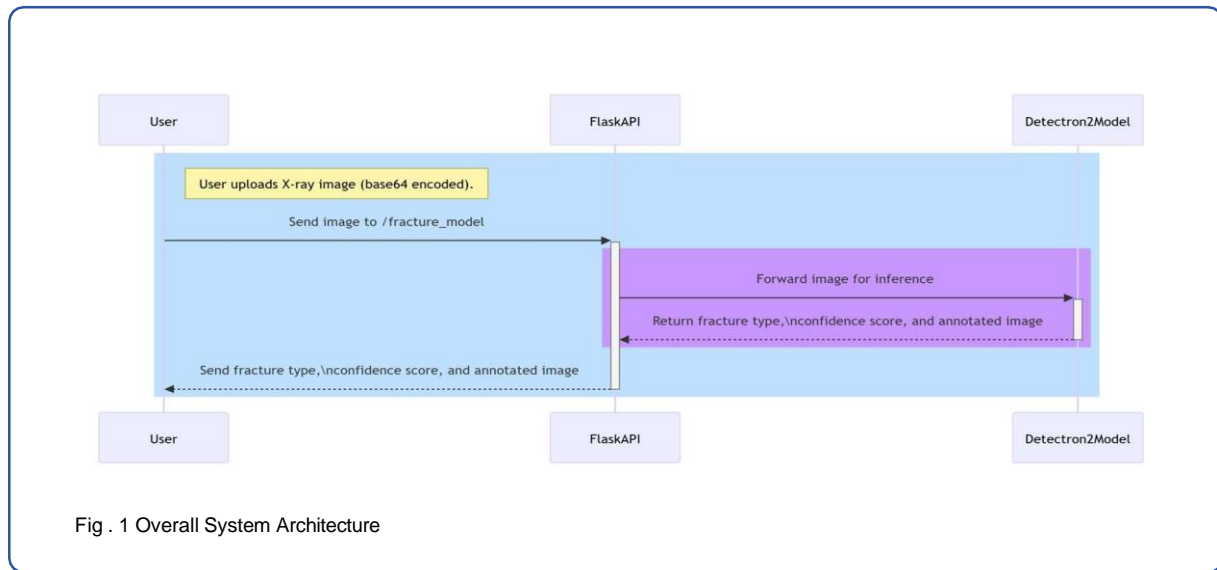
Reinforcement Learning: Think of it like training a dog—you reward or punish based on behavior. It learns by trial and error. Cool stuff, but again, not directly useful for fracture detection.

Now coming to my project, I trained a model using supervised learning with a set of X-ray images labeled according to fracture types like forearm fracture, shoulder fracture, etc.

This helps the system learn exactly what features indicate a fracture and what don't.

Also, here's a small reality check—most of the traditional machine learning methods assume that all your data fits neatly into memory. But X-ray datasets (especially annotated COCO-style datasets) can get massive. That's why I had to go with deep learning, which is way better at handling large-scale, complex visual data. And with tools like Detectron2, it becomes a lot easier to manage that kind of heavy lifting.

1.2.2 A Few Smart Learning Methods That Really Help Okay, so now that we've got the basics out of the way,



let's talk about some of the advanced learning techniques that actually helped make my bone fracture detection model better. These methods are not just about using one fixed algorithm— they're more like smart approaches that help machines understand complicated data (like X-rays) more effectively.

### 1. Representation Learning

When you're working with medical images, especially X-rays, the data isn't always easy to understand. You're dealing with all kinds of patterns, shadows, bones, and sometimes even blurry stuff. That's where representation learning comes in. The whole idea here is to let the model automatically learn the best features to look at—like what parts of the image matter the most when spotting a fracture. In my project, the Faster R-CNN model (from Detectron2) handles this by learning from thousands of image examples. It figures out what areas typically look like fractures and what's normal, without me having to manually define rules like “if a line is broken here, it's a fracture.” That saves a lot of time and also improves accuracy.

### 2. Deep Learning

Let's be real—deep learning is the superstar here. It's basically the reason why my fracture detection model can do what it does. Unlike older machine learning models that only look at surface-level features, deep learning models use layers of neural networks to learn from raw images in a much deeper way.

For my project, I used a deep learning model called Faster R-CNN, which not only detects if there's a fracture but also draws a box around it—so you get both detection and localization. That's a huge upgrade compared to old-school systems. And the best part? It

keeps getting better the more you train it.

### 3. Distributed and Parallel Learning

Now I didn't use distributed learning directly in my project (because I was running it on CPU, not on a fancy cluster), but it's still worth mentioning. When datasets get really big—like if a hospital is dealing with thousands of X-rays a day—it makes sense to split the work across multiple machines. That's what distributed learning does. It helps train models faster and on larger datasets by dividing the work.

If I ever scale this system up to something like a hospital network, distributed learning is gonna be the way to go.

### 4. Transfer Learning

Here's something I actually used: transfer learning. Basically, instead of training a model from scratch (which takes forever and needs tons of data), I used a model that was already trained on a large dataset (like COCO), and just fine-tuned it on my fracture data. It's like taking a student who already knows general anatomy and teaching them how to specifically detect fractures. Much faster, and still super accurate.

### 5. Active Learning

Okay, I didn't fully implement this, but it's something I want to explore later. In simple terms, active learning helps you deal with situations where you have tons of images, but only a few are labeled. The model learns from the few labeled ones and then asks you, “Hey, can you label this one? It looks important.” That way, you don't waste time labeling everything. Super useful when data labeling is slow (like in medical datasets).

### 6. Kernel-Based Learning

This one is more academic, and not really what I used,

but it's cool to mention. Kernel-based methods, like SVMs, help you deal with non-linear data—data that doesn't follow a simple line or curve. These methods work well when data is tough to separate. But honestly, for image-based stuff like fracture detection, deep learning beats kernel methods most of the time these days.

*1.3 Real-World Challenges I Faced While Building the Bone Fracture Detection System* Even though deep learning sounds super powerful—and it is—it's not like you plug in data and boom, you get perfect results. Especially in a medical use case like bone fracture detection, the challenges really stack up. It's not just about model accuracy. You also have to think about huge image sizes, odd data formats, system speed, label quality, and getting real value from the predictions. So in this section, I'm just gonna walk you through five major problems I personally ran into while working on this project, and how I tried to deal with them.

*1.3.1 Problem #1: Dealing with Massive X-ray Data (The "Volume" Problem)* One of the first things I realized is that medical data isn't light—literally. X-rays are huge, and when you have hundreds or thousands of them, it starts hitting your RAM, GPU, storage—everything. My laptop was definitely not happy.

How I handled it:

I used the COCO dataset format mainly because Detectron2 plays really well with it. It made organizing annotations and training data way easier. For the actual model training, I broke the data into smaller batches and used image resizing (but carefully, since resizing too much can erase small fracture details). For bigger projects, I'd definitely recommend using GPUs on Google Colab Pro, or even training on the cloud using something like AWS EC2 or Paperspace. Distributed training is a thing too, but I kept things single-machine for now.

*1.3.2 Problem #2: Every X-ray Is Different (The "Variety" Problem)* So here's the thing—just because all the data is "X-rays" doesn't mean they all look or behave the same. Some are frontal, some lateral. Some are bright, others super dark. Some images even had random text printed over them. This inconsistency really messes with the model's understanding of what a fracture looks like.

What I did:

I ended up writing a custom data preprocessor. It handled things like contrast normalization, orientation fixes, and

filtering out really bad samples. I also used image augmentations (flips, rotations, brightness changes) so the model wouldn't overfit to a specific style of X-ray. Detectron2 made that part easy since you can just plug augmentations right into the data loader. If I were working with multi-source hospital data, I'd consider using domain adaptation techniques or transfer learning to better handle the differences.

*1.3.3 Problem #3: Fast Inference Needed for Real Use (The "Velocity" Problem)* This was a big one. In real life, if a doctor wants to use your tool, they're not gonna wait 2 minutes for it to spit out a result. You need fast predictions. Especially in emergencies where fracture detection can't be delayed.

My solution:

I wrapped my trained model inside a Flask API, which loads the model once and keeps it in memory. That way, whenever a request comes in with a new X-ray, it returns the result in a few seconds—way faster than reloading the model every time. For future versions, I'm thinking of converting the PyTorch model to ONNX and using something like NVIDIA Triton or TensorRT for even faster inference.

*1.3.4 Problem #4: Messy or Missing Labels (The "Veracity" Problem)* Let's be real—getting X-ray images is one thing, getting high-quality labeled data is a whole different headache. A lot of open datasets either don't have precise bounding boxes or the annotations are vague like "fracture somewhere here," which isn't very helpful for training.

What I did:

I cleaned the dataset manually where I could, and skipped samples that had questionable labels. I also made sure that my loss functions weren't too sensitive to noise. In the future, I'd love to implement an active learning loop where the model flags "uncertain" predictions and a radiologist can confirm them. That way, the model keeps improving without needing tons of new labeled data every time.

*1.3.5 Problem #5: Finding Value in "Empty" Images (The "Value" Problem)* A huge part of this project is that not every X-ray actually has a fracture. So if your model sees 100 images and only 10 have fractures, you've got a major class imbalance. Plus, even when there is a fracture, it might be a tiny crack that's easy to miss—both for humans and models.



How I approached it:

That's where Faster R-CNN helped. Since it's a region-based detector, it doesn't try to classify the whole image—it focuses only on areas that might be interesting (like small cracks). I

also customized the anchor box sizes and ROI thresholds in Detectron2 so the model was better tuned to pick up on smaller features. Training with a mix of easy and hard examples helped a lot too.

### 1.3.6 Summary of My Experience

Each of these challenges was like a mini project of its own. Building a fracture detection system isn't just about deep learning—it's also about data wrangling, optimization, deployment, and a bit of problem-solving creativity. The key takeaway? You don't need fancy infrastructure or a PhD to do useful AI. But you do need to understand your data really well and be ready to adapt as new problems show up. Every time I fixed one thing, something else would break—but that's just how building real projects goes, right?

## 1.4 How Deep Learning in Medical Imaging Connects with Signal Processing Ideas

While I mostly worked with deep learning tools like Detectron2 and Faster R-CNN, I started realizing that a lot of what we're doing—especially when dealing with X-ray images—is basically an evolved form of signal processing. I mean, X-rays are just grayscale images, and image data is basically structured signals. So in a way, deep learning in vision is like modern signal processing, just way more data-hungry and powerful.

Let me break down how these two areas (ML and SP) kind of come together, especially in a project like mine.

### 1.4.1 Statistical Learning Meets Medical Imaging

When I started training my model, I had to think about things like how much data I have, what kind of patterns the model is learning, and whether it's overfitting or not. That's basically statistical learning in action.

What's cool is that X-ray images have certain statistical patterns. For example, fractures often show up as high-frequency edges or sudden discontinuities in the bone structure. Back in traditional SP, edge detection used to be done with Sobel filters or wavelets. Now, with CNNs, it's like those filters are being learned automatically. In my Faster R-CNN model, the backbone (ResNet50) is doing that kind of low-level pattern extraction without me hard-coding anything. That's the power of statistical learning meeting signal data.

### 1.4.2 Optimization Behind the Scenes

Training deep models like Faster R-CNN might look like magic, but under the hood, it's all just one big optimization problem. You've got loss functions trying to minimize classification errors, bounding box regression errors, and everything in between.

Here's where it overlaps with signal processing again. In SP, optimization methods like convex optimization or signal reconstruction via sparse coding were huge. These ideas are still at play—just scaled up and more automated. Even the Detectron2 framework uses optimizers like SGD and Adam, which are just iterative solvers working on huge multidimensional loss landscapes. So whether you're trying to restore a corrupted audio signal or detect a fracture in an X-ray, you're basically solving some variant of "What's the best match for this pattern?"

### 1.4.3 Real-Time Processing and Streamlined Inference

In signal processing, you often work with real-time data—like audio streams or live sensor readings—and your algorithm has to respond *now*, not later. That's super relevant for my use case too.

Let's say this fracture detection system is deployed in a clinic. A doctor uploads an X-ray and expects a result in seconds. That means I can't afford to re-load a 200MB model every time someone hits "submit." So I applied a bit of the SP mindset here: keep the model warm in memory (just like a buffer in signal systems) and run inference immediately, almost like streaming data.

If I were to push this further—say, on an embedded system or mobile device—I'd totally borrow ideas from SP again, like model pruning or quantization (basically compressing models like we used to compress signals).

### 1.4.4 Filtering Out the Noise

Another direct connection between SP and my project is noise reduction. Some X-ray images are noisy—blurry, low contrast, or scanned at a bad angle. And I noticed that my model's performance dropped on those. That's no different from signal denoising techniques used in audio or radar.

To fix that, I added basic image preprocessing steps like histogram equalization and contrast normalization. These are classic signal processing tricks—but they helped my deep learning model perform way better. If I had more time, I'd try integrating a dedicated image enhancement module

Table 1: Model Training Configuration Table

Hyperparameter	Value
Backbone	ResNet101 + FPN
Batch Size per Image	512
Epochs	1800
Score Threshold	0.3
Number of Classes	8
Device	GPU

before detection, like a small CNN trained to “clean up” X- rays.

#### 1.4.5 Bringing It All Together

So yeah, while I was knee-deep in coding and model training, I kept stumbling into ideas that came straight from the signal processing world—just with more compute and data behind them. Whether it’s filtering, optimization, real-time inference, or denoising, the overlap is very real.

Deep learning doesn’t replace SP—it extends it. And in projects like this, understanding both gives you a huge edge.

**1.5 What Still Needs Work – Open Problems and Future Ideas** So yeah, building my bone fracture detection system with deep learning was a big step—but if I’m being honest, we’re just scratching the surface here. As cool as Faster R-CNN and Detectron2 are, there’s still a ton of stuff that needs to be figured out before this becomes a truly reliable, real-world tool that doctors can trust every single time.

This section is more like a wishlist or a set of “here’s what we still need to figure out” notes from my end—stuff that I either didn’t get time to explore or that’s still a bit messy in the whole field of deep learning for medical imaging.

**1.5.1 The “Meaning” Problem: Not All X-rays Are Equal** One big issue is how different hospitals or datasets interpret the same thing. For example, a fracture in one dataset might mean something totally different in another because of how the radiologist labeled it. Some datasets are super strict with bounding boxes, others just draw a blob around a region.

What we need:

We really need standardization—or at least smart models that can adapt to slightly different labeling conventions. Transfer learning and domain adaptation are a good start, but there’s still a gap when it comes to *context-aware learning*. Maybe combining deep learning with some kind of rule-based system (like “if the fracture is less than 2 cm, use this criteria”) could help add more intelligence.

**1.5.2 Training Data vs Labeling Cost: The Classic Trade-off** Let’s be honest, getting a large dataset is hard enough, but getting it *properly labeled*? That’s even worse. You need expert radiologists for this stuff, not just anyone with a mouse. What we need: Smarter labeling techniques. I’d love to see active learning systems where the model asks humans only for the “hard” examples, instead of labeling everything manually. Also, self- supervised learning could be a game-changer here—let the model learn representations even from unlabeled X- rays and fine-tune it later on the labeled ones.

**1.5.3 Mixing It All Together: Deep Learning + Other Tech** Right now, most people (including me) throw a deep model at the problem and try to make it work. But what if we combine deep learning with other techniques like medical ontologies, cloud-based processing, or even classical image processing pipelines?

What we need:

Hybrid systems. Imagine a pipeline that starts with rule-based filtering, then uses deep learning for detection, and finally adds a cloud-based second-opinion module. Sounds complicated, but if done right, this could be super powerful. It’s like building a full-stack AI for medicine, not just one smart model.

**1.5.4 Privacy and Ethics: The Stuff We Can’t Ignore** Working with medical data isn’t just about accuracy—it’s about trust. If hospitals are going to use tools like this, they need to know that patient privacy is protected. And honestly, I didn’t dive deep into this yet, but it’s definitely something that matters a lot.

What we need:

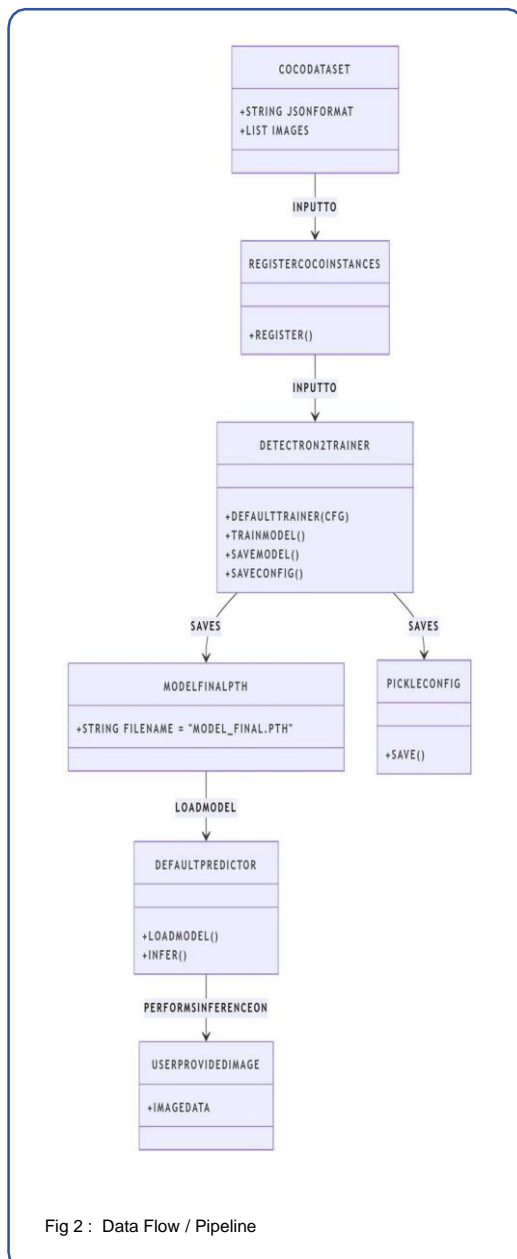
Better ways to anonymize medical images and still retain their usefulness. Also, federated learning sounds like a cool

idea— train your model across hospitals without actually moving the data around. That way, privacy is preserved, but the model still learns from a wider pool.

### 1.5.5 Taking It from Demo to Deployment

It's one thing to get 90% accuracy on your test set. It's a totally different thing to make your system usable in a hospital, by actual doctors, on random new images from different machines.

What we need:



A better bridge between research and deployment. Right now, most models (including mine) are stuck as local prototypes or Flask apps. The next step is to turn them into production-grade

services—with monitoring, version control, fail-safes, and

all that serious stuff. Tools like ONNX, Docker, and Kubernetes are probably going to be part of that path.

Final Thought on This Section

So yeah, while I've made good progress, there's still a long road ahead before deep learning can be fully trusted in something as sensitive as healthcare. But that's what makes it exciting—there's so much room to explore, improve, and innovate. Whether it's smarter training methods, better privacy protection, or just making things more usable in the real world, there's a lot we can still build.

And who knows? Maybe this fracture detection project is just the beginning

## 2. CONCLUSIONS

So yeah, to sum everything up—working on this bone fracture detection system has been a wild ride. I started off just wanting to build a deep learning project with some real-world value, and I ended up diving deep into medical imaging, COCO annotations, Faster R-CNN models, Detectron2, and wrapping it all with a Flask API.

But along the way, I realized something important: training a model is just one part of the story. Once you bring in real X- rays, huge file sizes, incomplete annotations, deployment needs, and even ethical stuff like privacy— everything suddenly becomes a lot more complex, and honestly, more interesting. This project taught me how deep learning isn't just about high accuracy—it's about understanding the problem, cleaning your data, designing the right pipeline, and building something that people might actually use. And that's what I loved about it. It wasn't perfect, but it was real. Of course, there's still plenty of room to improve—from smarter labeling methods and better generalization to real-time performance and scaling this for clinical use. But this work shows that even with limited resources, it's totally possible to build a meaningful AI system that could assist in healthcare someday. So, in the end, this project wasn't just about detecting fractures. It was about learning how to think like a builder, a problem solver, and maybe even a researcher. And I'm excited to keep going.

So yeah, to sum everything up—this bone fracture detection system ended up being way more than just a deep learning project. What started off as an idea to “just build something cool with AI” turned into this deep dive into the whole ecosystem of medical imaging, annotations, model optimization, and deployment. I thought I was just going to train a model, but I quickly learned that's only one chapter in a much bigger story.

When I started working with X-rays, I had no idea how messy and complex real-world data could be. Some images were too large to even load without crashing, annotations

weren't always clean or complete, and organizing everything into COCO format took way more effort than I expected. But figuring out how to deal with all that? That's what made this project feel real.

I used Detectron2 with Faster R-CNN because it's pretty robust and well-documented, especially for object detection tasks. And honestly, just getting it to detect different types of fractures—elbow, wrist, humerus, etc.—was such a satisfying milestone. But the more I worked on it, the more I realized how important the backend setup was. Wrapping it all up with a Flask API taught me a lot about how to actually serve a model, not just train one. It made the whole thing feel more like an actual product, something you could deploy and use in a real-world workflow.

And then there's the ethical and practical stuff that I didn't fully expect. Things like patient privacy, handling sensitive data, thinking about how a model like this could actually be used in clinics—not just as a demo, but as a decision support tool. That part really shifted my mindset. This wasn't just an academic project anymore—it was something that, in the future, could genuinely help people if built right.

The truth is, I ran into a bunch of challenges—resource constraints, no GPU, long training times, model generalization issues—but I realized that this is part of the game. It's not about having everything perfect. It's about figuring things out along the way, being resourceful, making the best of what you have, and still building something that works.

What I loved most was how this project pushed me to think beyond just accuracy and architecture. It made me think about the full pipeline—from dataset quality and model training to real-time inference and user interaction. And maybe more importantly, it made me think like a builder. Someone who doesn't just train models, but solves problems.

There's still a lot to improve—like optimizing it for better performance, using more diverse data for better generalization, exploring active learning for smarter labeling, and maybe even building a proper UI for doctors. But now I've got a foundation to work on. I've seen how a project can go from code to something that feels impactful. And that's what excites me.

So yeah, this wasn't just about detecting bone fractures. It was about learning how to build something real. And it's only the beginning.

## REFERENCES

- [1] Ren, S., He, K., Girshick, R., & Sun, J. (2015). *Faster R-CNN: Towards Real-Time Object Detection with Region Proposal Networks*. Advances in Neural Information Processing Systems, 28. [Link](#)
- [2] Girshick, R. (2015). *Fast R-CNN*. Proceedings of the IEEE International Conference on Computer Vision (ICCV), 1440–1448. [Link](#)
- [3] Redmon, J., & Farhadi, A. (2018). *YOLOv3: An Incremental Improvement*. arXiv preprint arXiv:1804.02767. [Link](#)
- [4] He, K., Gkioxari, G., Dollár, P., & Girshick, R. (2017). *Mask R-CNN*. Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2961–2969.
- [5] Lin, T.-Y., Maire, M., Belongie, S., et al. (2014). *Microsoft COCO: Common Objects in Context*. European Conference on Computer Vision (ECCV), 740–755. [Link](#)
- [6] Krizhevsky, A., Sutskever, I., & Hinton, G. E. (2012). *ImageNet Classification with Deep Convolutional Neural Networks*. Advances in Neural Information Processing Systems, 25. [Link](#)
- [7] Abadi, M., Barham, P., Chen, J., et al. (2016). *TensorFlow: A System for Large-Scale Machine Learning*. Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation (OSDI), 265–283. [Link](#)
- [8] Paszke, A., Gross, S., Massa, F., et al. (2019). *PyTorch: An Imperative Style, High-Performance Deep Learning Library*. Advances in Neural Information Processing Systems, 32.
- [9] Facebook AI Research (FAIR). *Detectron2: A PyTorch-based modular object detection library*. [GitHub Repository](#)
- [10] Flask Documentation. *Flask: Web development, one drop at a time*. Official Documentation
- [11] Howard, A. G., Zhu, M., Chen, B., et al. (2017). *MobileNets: Efficient Convolutional Neural Networks for Mobile Vision Applications*. arXiv preprint arXiv:1704.04861.
- [12] Tan, M., & Le, Q. V. (2019). *EfficientNet: Rethinking Model Scaling for Convolutional Neural Networks*. Proceedings of the 36th International Conference on Machine Learning (ICML), 6105–6114.
- [13] He, K., Zhang, X., Ren, S., & Sun, J. (2016). *Deep Residual Learning for Image Recognition*. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 770–778.
- [14] Simonyan, K., & Zisserman, A. (2015). *Very Deep Convolutional Networks for Large-Scale Image Recognition*. International Conference on Learning Representations (ICLR). [Link](#)
- [15] Szegedy, C., Liu, W., Jia, Y., et al. (2015). *Going Deeper with Convolutions*. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 1–9.
- [16] Russakovsky, O., Deng, J., Su, H., et al. (2015). *ImageNet Large Scale Visual Recognition Challenge*. International Journal of Computer Vision, 115(3), 211–252.
- [17] Deng, J., Dong, W., Socher, R., et al. (2009). *ImageNet: A Large-Scale Hierarchical Image Database*. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 248–255.
- [18] Zhou, B., Khosla, A., Lapedriza, A., et al. (2016). *Learning Deep Features for Discriminative Localization*. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 2921–2929.
- [19] Selvaraju, R. R., Cogswell, M., Das, A., et al. (2017). *Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization*. Proceedings of the IEEE International Conference on Computer Vision (ICCV), 618–626.
- [20] Dosovitskiy, A., Beyer, L., Kolesnikov, A., et al. (2021). *An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale*. International Conference on Learning Representations (ICLR).
- [21] Vaswani, A., Shazeer, N., Parmar, N., et al. (2017). *Attention is All You Need*. Advances in Neural Information Processing Systems, 30.
- [22] Carion, N., Massa, F., Synnaeve, G., et al. (2020). *End-to-End Object*



- Detection with Transformers*. European Conference on Computer Vision (ECCV), 213–229.
- [23] Chen, T., Kornblith, S., Norouzi, M., & Hinton, G. (2020). *A Simple Framework for Contrastive Learning of Visual Representations*. International Conference on Machine Learning (ICML), 1597–1607.
- [24] Grill, J.-B., Strub, F., Altché, F., et al. (2020). *Bootstrap Your Own Latent: A New Approach to Self-Supervised Learning*. Advances in Neural Information Processing Systems, 33, 21271–21284.
- [25] Chen, X., Fan, H., Girshick, R., & He, K. (2020). *Improved Baselines with Momentum Contrastive Learning*. arXiv preprint arXiv:2003.04297.
- [26] Gulrajani, I., Ahmed, F., Arjovsky, M., et al. (2017). *Improved Training of Wasserstein GANs*. Advances in Neural Information Processing Systems, 30.
- [27] Goodfellow, I., Pouget-Abadie, J., Mirza, M., et al. (2014). *Generative Adversarial Nets*. Advances in Neural Information Processing Systems, 27.
- [28] Radford, A., Metz, L., & Chintala, S. (2016). *Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks*. International Conference on Learning Representations (ICLR).
- [29] Isola, P., Zhu, J.-Y., Zhou, T., & Efros, A. A. (2017). *Image-to-Image Translation with Conditional Adversarial Networks*. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 1125–1134.
- [30] Zhu, J.-Y., Park, T., Isola, P., & Efros, A. A. (2017). *Unpaired Image-to-Image Translation using Cycle-Consistent Adversarial Networks*. Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2223–2232.
- [31] Huang, G., Liu, Z., Van Der Maaten, L., & Weinberger, K. Q. (2017). *Densely Connected Convolutional Networks*. Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR), 4700–4708.
- [32] Ioffe, S., & Szegedy, C. (2015). *Batch Normalization: Accelerating Deep Network Training by Reducing Internal Covariate Shift*. International Conference on Machine Learning (ICML), 448–456.
- [33] Srivastava, N., Hinton, G., Krizhevsky, A., et al. (2014). *Dropout: A Simple Way to Prevent Neural Networks from Overfitting*. Journal of Machine Learning Research, 15(1), 1929–1958.
- [34] Kingma, D. P., & Ba, J. (2015). *Adam: A Method for Stochastic Optimization*. International Conference on Learning Representations
- [35] (ICLR).
- [36] Zeiler, M. D., & Fergus, R. (2014). *Visualizing and Understanding Convolutional Networks*. European Conference on Computer Vision (ECCV), 818–833.
- [37] Yosinski, J., Clune, J., Bengio, Y., & Lipson, H. (2014). *How Transferable Are Features in Deep Neural Networks?* Advances in Neural Information Processing Systems, 27.
- [38] Pan, S. J., & Yang, Q. (2010). *A Survey on Transfer Learning*. IEEE Transactions on Knowledge and Data Engineering, 22(10), 1345–1359.
- [39] Shorten, C., & Khoshgoftaar, T. M. (2019). *A Survey on Image Data Augmentation for Deep Learning*. Journal of Big Data, 6(1), 60.
- [40] Zhang, H., Cisse, M., Dauphin, Y. N., & Lopez-Paz, D. (2018). *mixup: Beyond Empirical Risk Minimization*. International Conference on Learning Representations (ICLR).
- [41] Cubuk, E. D., Zoph, B., Mane, D., et al. (2019). *AutoAugment: Learning Augmentation Policies from Data*. Proceedings of the IEEE/C