Railway Track Crack Detection Using Sound Signal Analysis and Gen AI Models

AUTHORS

1. **K.Kirubagaran,**
Pg & Research Department of Computer Science,
Sri Ramakrishna College Of Arts & Science,
Kirubagaran460@gmail.com

1. **DR R Nagarajan,**
Assistant Professor,
Pg & Research Department of Computer Science,
Sri Ramakrishna College Of Arts & Science,
rnagarajan.snr@gmail.com

ABSTRACT

Keeping railway tracks safe really depends on spotting cracks, corrugation, and weld failures as soon as possible. The old-school ways—manual inspections, ultrasonic testing vehicles, and regular maintenance—just take too long, cost too much, and can't watch the tracks nonstop.

So, in this work, the goal is to build a smarter Railway Track Crack Detection System that listens for trouble. The idea is to grab the sounds that happen when train wheels meet the rails, then dig into those sounds using signal processing techniques like Mel-Frequency Cepstral Coefficients, Spectral Feature Extraction, and Fractal Dimension Analysis. These methods help pull out details from the noise that can reveal if something's wrong.

But there's a catch: there aren't enough examples of damaged tracks to train a good model. To get around that, we use Generative AI to create new, realistic defect samples and balance the dataset.

Once we've got solid data, we turn to machine learning—Random Forest, Support Vector Machine, XGBoost, and some ensemble methods—to teach the system how to tell healthy tracks from damaged ones. The end goal? Faster, smarter, and more reliable railway safety.

Keywords: Railway Crack Detection, Acoustic Signal Processing, Generative AI, MFCC, Spectral Analysis, Ensemble Learning, Infrastructure Monitoring.

INTRODUCTION

Trains haul millions of people and tons of cargo every single day, so the strength of railway tracks isn't just important—it's crucial. When cracks go unnoticed or the rails start wearing out, things can fall apart fast. We're not just talking about money here; people's lives are at stake. Tracks put up with a lot. Picture it: massive trains pounding over them, wild swings in temperature, rain, snow, heat, you name it. All that punishment adds up, and little cracks start forming. At first, you probably wouldn't even spot them, but give it time and those cracks can grow into real problems.

Rail companies have usually relied on ultrasonic scanners or teams of workers for visual checks. These methods get the job done, but they've got some big drawbacks. They're done on a schedule, not around the clock, so it's easy for something to slip by between inspections. Plus, it costs a lot. You need trained people, special gear, and it's nearly impossible to cover every stretch of track, especially with huge rail networks. The old ways have kept trains moving for decades, but let's be honest—they're starting to fall short for what we need today.

Signal processing and AI are shaking up how we find defects—no need to tear everything apart anymore. Picture a train rolling over a bad patch of track. The noise changes. It's subtle, but it's there. The rhythm, the pitch, the way the sound lingers—all those little details give away what's happening under the wheels. With the right tech, you can grab those clues and figure out what's wrong, fast. That's really what this research is all about. We're making a smart system that listens to the tracks and uses Generative AI to catch cracks before they get serious. It's a shift from the old, hands-on inspections to something automatic and built on real data—something that actually scales. In the end, this approach helps us catch issues early and do it with real confidence.

OBJECTIVE OF THE STUDY

We're aiming to build a smarter way to find cracks in railway tracks by listening to the sounds they make and letting generative AI figure out what's a defect and what's not. Here's how it goes: first, we set up a system that captures the sounds from the tracks, gets rid of the background noise, and focuses on the parts of the audio that actually tell you something about the track's health. We use things like MFCC extraction, spectral centroid and rolloff, spectral bandwidth, spectral flux, zero-crossing rate, and fractal dimension analysis—a whole toolkit designed to dig into the weird details and patterns in those sounds. The goal? Spot problems early, fix them fast, and keep the trains running safely.

This study wants to solve a big issue: there just isn't enough real-world data on railway defects. So, the team steps in with Generative AI and cooks up lifelike defect sound samples. These fresh samples make training way more effective and help even out the class imbalance that usually throws off accuracy. They also run tests with different machine learning models—Random Forest, XGBoost, Support Vector Machines, and ensemble voting—to see which one actually gets the job done. The goal? Build a system that's both scalable and budget-friendly, one that monitors railways in real time and helps keep trains safer and more reliable through predictive maintenance.

METHODOLOGY

This study rolls up its sleeves and gets right into the nitty-gritty of engineering, working through everything from grabbing the data to actually building and launching the system. It starts with collecting railway acoustic signals—sometimes we pull them straight from the tracks, other times we use models to whip up synthetic samples that sound like either healthy or damaged rails. Healthy rails usually give off these deep, steady rumbles and regular clunks from the wheels. But when there's something wrong, like cracks or bad welds, you hear sharper, messier sounds—lots of sudden pops and odd patterns that just don't fit.

After we record the audio, we clean it up. That means we set the volume levels, wipe out background noise, and make sure every sample runs at the same sampling rate. Then we jump into feature extraction. We tear into the audio with advanced spectral analysis. MFCC coefficients help us catch the kind of details people actually hear. Other features—like spectral centroid, bandwidth, rolloff, and spectral flux—help us chart how the energy in the sound spreads out and moves. We also check out things like zero-crossing rate and spectral kurtosis to spot those sharp, weird signals, and we throw in fractal dimension analysis to see just how complicated the waveforms get.

We deal with dataset imbalance by turning to generative AI, creating extra examples of defective sounds so our system has more to learn from. Diffusion-based modeling lets us simulate crack signatures that sound real enough to fool even careful listeners. This makes our classifier sharper during training. For the actual machine learning, we use models like Random Forest, XGBoost, and Support Vector Machines. We train them on feature vectors straight from the data, then blend their predictions with an ensemble method. That move gives us more reliable results and keeps false alarms low. And the best part? Everything runs through a simple web interface—just upload a railway sound recording and get instant feedback about cracks.

SYSTEMS

We ran system tests to see if everything actually worked together like it should in real life. First, we checked each module by

itself. We needed to know if the system really caught things like MFCC coefficients, spectral features, and fractal measures from the audio. We even listened to the synthetic samples it made, just to make sure the new data actually sounded real, not obviously fake.

Then we moved on to the classifier. We split the data into training and testing sets, and checked accuracy, precision, recall, and F1-score. The confusion matrix made it easy to spot mistakes—especially false positives and false negatives. That kind of error matters. If the system misses a crack in railway tracks, that's just not acceptable.

We didn't stop there. We pushed the system, running batch processing simulations to see how it managed a bunch of samples at the same time. We wanted to see if it stayed fast and reliable, even when the workload got heavy. On top of that, we checked the user experience. People needed to upload audio, see the features, and understand the results without running into any roadblocks. Last thing—we double-checked security. Every file upload was handled safely, so nothing could mess with the system's integrity.

RESULTS

Testing the new railway crack detection system turned out better than we expected. Mixing acoustic signal processing with Generative AI gave the system a real edge—it caught defects a lot more accurately. Once we started digging into the features, it got pretty clear: cracked rails showed higher spectral flux, bigger swings in spectral centroid, and way more fractal complexity compared to healthy rails.

Throwing in synthetic defect samples changed the game. The classifier handled everything with more confidence, didn't overfit, and the sample set felt more balanced overall. Out of all the models we tried, the ensemble method—using both Random Forest and XGBoost—really delivered. It hit the sweet spot for accuracy and stayed stable, plus it kept precision and recall high. In plain terms, it picked up the bad rail sections and barely sounded any false alarms.

We timed it, too. The model processed uploaded audio in just a few seconds, so you're basically looking at near real-time monitoring—no waiting around. Bottom line, this isn't just a cool research project. Combining acoustic signals with Generative AI actually works in the field and can scale up fast, making it a solid, smarter alternative to traditional inspection.

CONCLUSION

The Railway Track Crack Detection System changes the game for how we keep an eye on railway tracks. Instead of tearing up

rails or blowing the budget, this system listens for trouble. It grabs sounds from the tracks, pulls out the important details, and uses machine learning to figure out if there's a crack. Generative AI steps in and makes the whole setup even tougher—so it works well, even when there aren't tons of examples of broken tracks to learn from.

Put it all together, and you get a smart mix of signal processing, AI, and generative models working for railway safety. And honestly, it's not just about finding cracks. This approach opens the door to predictive maintenance, smarter IoT monitoring, and automated health checks. The research backs it up: smart acoustic monitoring can stop accidents before they happen, keep the rails in better shape, and make train travel safer for everyone.

BIBLIOGRAPHY

[1] S. Mallat, *A Wavelet Tour of Signal Processing*, 3rd ed. Burlington, MA, USA: Academic Press, 2009.

[2] L. Rabiner and B. H. Juang, *Fundamentals of Speech Recognition*. Englewood Cliffs, NJ, USA: Prentice Hall, 1993.

[3] S. Theodoridis and K. Koutroumbas, *Pattern Recognition*, 4th ed. Burlington, MA, USA: Academic Press, 2009.

[4] T. Li, M. Chen, and X. Zhang, "Railway track crack detection based on acoustic emission signals," *IEEE Sensors Journal*, vol. 19, no. 15, pp. 6097–6105, 2019.

[5] J. Lin, Q. Wang, and Y. Li, "Automatic rail defect detection using audio signal processing and machine learning," *Mechanical Systems and Signal Processing*, vol. 120, pp. 328–345, 2019.

[6] D. P. Kingma and M. Welling, "Auto-Encoding Variational Bayes," in *Proc. International Conference on Learning Representations (ICLR)*, 2014.

[7] I. Goodfellow et al., "Generative Adversarial Networks," in *Proc. Advances in Neural Information Processing Systems (NeurIPS)*, 2014.

[8] J. Ho, A. Jain, and P. Abbeel, "Denoising Diffusion Probabilistic Models," in *Proc. Neural Information Processing Systems (NeurIPS)*, 2020.

[9] B. Logan, "Mel Frequency Cepstral Coefficients for Music Modeling," in *Proc. International Symposium on Music Information Retrieval (ISMIR)*, 2000.

[10] M. Müller, *Fundamentals of Music Processing: Audio, Analysis, Algorithms, Applications*. Cham, Switzerland: Springer, 2015.