The Thirty-Ninth AAAI Conference on Artificial Intelligence (AAAI-25)

Hyper-Spectral Image Generation from Frequency Spectrums

Ruth-Emely Pierau[1,2]

[1] Department of Data Science and Artificial Intelligence, Monash University
[2]Future Fibre Technologies, 10 Hartnett Cl, Mulgrave VIC 3170, Australia

Introduction

My thesis primarily focuses on hyper-spectral image generation from frequency spectrums for downstream computer vision tasks. Hyper-spectral images are images with more than three channels commonly created by special hyper-spectral cameras or from frequency spectrums of various sensing applications such as radargrams or distributed acoustic sensing (DAS) systems. The range of frequencies considered in a frequency spectrum is typically too large to map one frequency to one image channel, i.e. we generally consider a frequency spectrum of 2500 Hz. Frequencies need to be binned together in frequency bands where each band forms one image channel. Usually, frequency bands are created either by expert knowledge or trial-and-error (Stork et al. 2020; Liu et al. 2019).

I research how filters can be trained to automatically select frequencies and bin them into frequency bands. My aim is to represent a variety of signal information and decrease noise. Signal representation is optimized for object detection on time-sequenced images with a set number of image channels. The object detection task consists of localizing and classifying events in the generated hyper-spectral images. Events are typically types of intrusions, structural changes, or defined actions and structures, e.g. climbing of a fence. Events and noise often share at least some frequencies and vary between application types. My research question is how hyper-spectral images can be automatically constructed from frequency spectrums for downstream computer vision tasks.

In my research, I mainly work with commercially deployed DAS systems from my industry partner, Future Fibre Technology, an AVA Risk Group company. DAS systems are capable of continuously monitoring vibrations and acoustic signals over long distances in real-time, with a high sensitivity and accuracy (Wang et al. 2014; Gabai and Eyal 2016). This has made them a valuable technology for the monitoring and protection of commercial assets and infrastructure. These include airports, power plants, and pipelines (Zhu et al. 2022; Chen et al. 2023).

A DAS system consists of a fiber optical cable and a controller unit. A laser sends pulses through the fiber. Scattered light is reflected back and sampled every 0.5 m. The pattern of the backscatter is changed by disturbances. Disturbances can be either event types or noise such as wind and wild animals. To improve classification, the DAS system is integrated with an object detector (Sha et al. 2021; Ma et al. 2023). The raw acoustic signal sensed by the fiber cable is transformed via Fast Fourier Transformation to a frequency spectrum which I use in my research to generate hyper-spectral images that represent distance in width, time in height, and frequencies in channel dimension.

Progress

I have progressed in my research to reliably and automatically create hyper-spectral images from frequency spectrums of DAS systems that outperform both the expert knowledge from my industry partner and baseline methods. We design filters to not only denoise the power spectrum but also select frequencies that represent the event information best. Previous approaches have not utilized the detection model to adapt the filtering of the power spectrum. Considering that the filtered power spectrum is the input for the detection model, the filters should be trained in conjunction with the neural network, not be modeled separately. The detection model serves as both measure and selector.

I developed a greedy selection method for feature forward selection that integrates the learning of the filters with the training of the object detector. The algorithm bins and averages over intervals of frequencies before mapping each averaged intensity to an image channel. Starting with a uniform frequency band configuration, each frequency band is individually evaluated, and the best performing frequency bands are selected to be halved. The other bands are dropped. For example, for a 10 channel image generated from a frequency spectrum from 0 Hz-2500 Hz, equal, non-overlapping frequency bands of 250 Hz are created, such as 0 Hz-250 Hz and 250 Hz-500 Hz. Each is mapped into an image channel. The frequency spectrum is thus split into 10 tuples each marking the start and end of a frequency band.

The image dataset is generated with the uniform filter configuration. After the detection model is trained on the generated images, each image channel is evaluated by the detection model and the five channels with the highest metric are selected and split; the others are dropped. This is repeated until our chosen metric no longer increases. Greedy selection focuses the image generation on the frequency bands that contain relevant event information and ensures that foreground objects are more strongly represented than noise.

I also modeled a greedy set search method for feature backward selection. The algorithm works similar to greedy selection but instead of evaluating each frequency band individually, a set of frequency bands is evaluated while one band is dropped. I have adapted greedy set search further to stop before half the channels are dropped if the evaluation metric falls below a set threshold. While greedy selection and greedy set search outperform the current state-of-the-art, both require a training time over several days.

I currently investigate models where the learnable filters are trained together with the object detector in several hours. I designed a simulated annealing based approach that changes the filter configuration every 10 to 20 epochs during model training and evaluates whether the new filters outperform the previous ones. New filter configurations are created by randomly dropping, splitting and merging frequency bands from the previous configuration. Model performance is better than the baseline methods but unable to outperform my greedy methods.

Anticipated Progress

In the future, I plan to investigate genetic algorithms to compare with my simulated annealing based approach. I hope that by replacing simulated annealing with a genetic algorithm, the new filter configuration inherits well working features from its parent configuration. Instead of limiting the child filter configuration to one change per mutation, the crossover of two parents should introduce a wider variety of potential frequency bands while containing randomness in the search space to an elite population.

I also will transfer my methods from frequency spectrums obtained from a DAS system to radargrams obtained by a probe. The specifications of radargrams differ from those of a DAS system and they encounter different challenges. Adapting and transferring my models will add more depth and comprehension to the construction of hyper-spectral images from frequency spectrums.

My anticipated progress aligns with my proposed research plan to firstly investigate the generation of hyper-spectral images from frequency spectrums of a DAS system for object detection tasks; secondly to further examine and optimize hyper-spectral image generation methods under different constraints such as real-time evaluation, feasible training time for deployed systems, and full automation; and thirdly to transfer models between different application and installations.

Anticipated Thesis Contributions

My anticipated thesis contributions are three technical chapters in total introducing five new methods that can be fully integrated into any DAS system with an object detector. I further expect that my model is fully capable of automatically learning filters for any frequency spectrum of either a DAS system or a radargram. I will present a method that transfers filters between different DAS system installations and between different radargrams. I have submitted one paper, and expect to submit at least one more paper following a patent application where I am a co-author. I plan to contribute part of my model implementations to the mmDetection framework, a computer vision library available on Github.

Timeline

I am currently in my last year of my PhD program. In the last 2.75 years, I have completed two research objectives: I am able to generate hyper-spectral images from frequency spectrums automatically and under different constraints. In the last year, I will develop a genetic algorithm and make my methods fully transferable between different applications. Lastly, I will write up my thesis.

Acknowledgments

This research was fully funded by the Australian Government through the Australian Research Council Industrial Transformation Training Centre in Optimisation Technologies, Integrated Methodologies, and Applications (OPTIMA), Project ID IC200100009, and by Future Fibre Technologies, an AVA Risk Group company. I would like to acknowledge my supervisors who have supported my research: Alaster Meehan, Hamid Rezatofighi, and Peter J. Stuckey.

References

Chen, S.; Han, J.; Sui, Q.; Zhu, K.; Lu, C.; and Li, Z. 2023. Advanced Signal Processing in Distributed Acoustic Sensors Based on Submarine Cables for Seismology Applications. *Journal of Lightwave Technology*.

Gabai, H.; and Eyal, A. 2016. On the sensitivity of distributed acoustic sensing. *Optics letters*, 41(24): 5648–5651.

Liu, H.; Ma, J.; Xu, T.; Yan, W.; Ma, L.; and Zhang, X. 2019. Vehicle detection and classification using distributed fiber optic acoustic sensing. *IEEE Transactions on Vehicular Technology*, 69(2): 1363–1374.

Ma, B.; Gao, R.; Zhang, J.; and Zhu, X. 2023. A YOLOX-Based Automatic Monitoring Approach of Broken Wires in Prestressed Concrete Cylinder Pipe Using Fiber-Optic Distributed Acoustic Sensors. *Sensors*, 23(4): 2090.

Sha, Z.; Feng, H.; Rui, X.; and Zeng, Z. 2021. PIG Tracking utilizing fiber optic distributed vibration sensor and YOLO. *Journal of Lightwave Technology*, 39(13): 4535–4541.

Stork, A. L.; Baird, A. F.; Horne, S. A.; Naldrett, G.; Lapins, S.; Kendall, J.-M.; Wookey, J.; Verdon, J. P.; Clarke, A.; and Williams, A. 2020. Application of machine learning to microseismic event detection in distributed acoustic sensing data. *Geophysics*, 85(5): KS149–KS160.

Wang, Z.; Zeng, J.; Li, J.; Peng, F.; Zhang, L.; Zhou, Y.; Wu, H.; and Rao, Y. 2014. 175km phase-sensitive OTDR with hybrid distributed amplification. In *23rd international conference on optical fibre sensors*, volume 9157, 1562–1565. SPIE.

Zhu, H.-H.; Liu, W.; Wang, T.; Su, J.-W.; and Shi, B. 2022. Distributed acoustic sensing for monitoring linear infrastructures: Current status and trends. *Sensors*, 22(19): 7550.