Segmentation of Mechanical Coupling Using Motion Amplification of Sub-Pixel Variation

Yi Cheng Zhu, Cheng Xu

CSC2529

Motivation

This research project presents a novel method to segment a video into mechanically coupled components based on the temporal frequency of the spatial phase variations. In industrial environments, such vibrations can be caused by misaligned bearings and other mechanical issues, and can lead to long-term failures if not caught early. The current standard method of detection involves the use of accelerometers, but it is not practical to deploy such sensors all over the infrastructures.

Our proposed method is based on local phase-based Eulerian motion amplification, which allows us to diagnose these vibrations by analyzing subtle sub-pixel intensity variations that are captured in video footage. By segmenting the video into mechanically coupled components, we can identify the source of the vibrations and prevent long-term failures such as loose fasteners.

Related Work

Previous studies, such as those by Fleet and Jepson [1990] and Wadhwa et al. [2013, 2014], have shown that local phase information can be used to extract sub-pixel-level information and create component image velocity fields. These works have implemented Eulerian motion amplification techniques using complex-valued steerable pyramids and Riesz Pyramids to amplify the space-domain phases of each pixel. These approaches have been successful in amplifying the motions to the point where they are detectable to the human eye.

However, these methods still require manual processing by experts to isolate undesired motions. Our proposed method aims to provide an additional layer of automation by clustering and segmenting the different types of vibrations based on their phase and amplitude. This will allow for more efficient and effective detection of vibrations in industrial settings.

References

David J Fleet and Allan D Jepson. Computation of component image velocity from local phase information. International journal of computer vision, 5(1):77–104, 1990.

Neal Wadhwa, Michael Rubinstein, Frédo Durand, and William T Freeman. Phase-based video motion processing. ACM Transactions on Graphics (TOG), 32(4):1–10, 2013.

Neal Wadhwa, Michael Rubinstein, Frédo Durand, and William T Freeman. Riesz pyramids for fast phase-based video magnification. In 2014 IEEE International Conference on Computational Photography (ICCP), pages 1–10. IEEE, 2014.

Reference: J. Immerkær, "Fast Noise Variance Estimation", Computer Vision and Image Understanding, Vol. 64, No. 2, pp. 300-302, Sep. 1996 [PDF]

New Technique

We adapted the pipeline pioneered by Wadhwa et al. for motion magnification. We first construct a Riesz pyramid for each frame of the video. We then extract the local amplitude and phase and a butterworth filter is applied between frames to extracted the wanted frequencies. At this point we collapse the Reisz pyramid prematurely and fit our segmentation on this embedding.

Motion Extraction with Riesz Pyramid

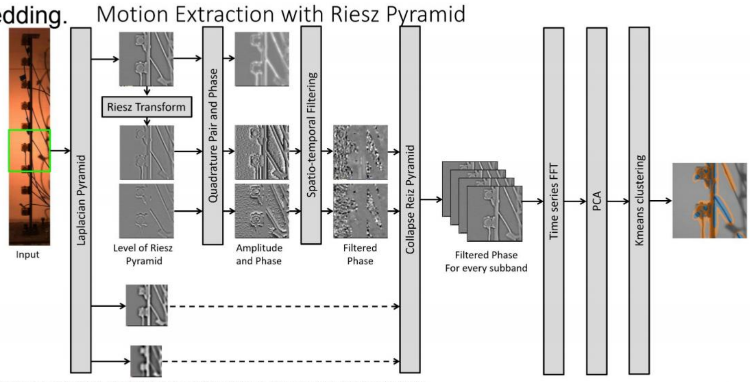

Neal Wadhwa, Michael Rubinstein, Frédo Durand, and William T Freeman. Riesz pyramids for fast phase-based video magnification. In 2014 IEEE International Conference on Computational Photography (ICCP), pages 1–10. IEEE, 2014.

For baseline comparisons, we elected two naive methods. We first average spatially to account for sensor and quantization error. We then feed the frequency domain stack into our standard segmentation pipeline. Another approach we explored is to extract out the top three frequency where the peaks were greatest and fit our Kmeans on the frequency informations.

Experimental Results

In Conclusion, our projects are extremely successful. The reisz pyramid based Segmentation behaves far superior. Not only are we able to distinguish the different frequencies, as our guitar examples show, but we are also able to segment based on the directionality of the vibration vectors, as our drum examples shows.

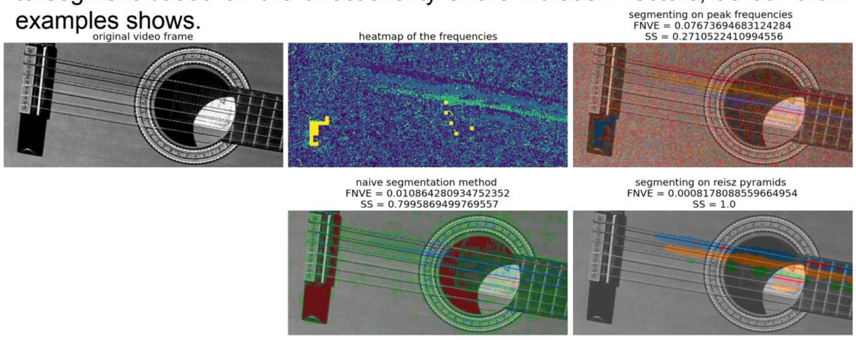

Since it is impossible to calculate PSNR ratio without a ground truth, we had to resort to Fast Noise Variance Estimation for noise metrics. On top of that, we used structural similarities to compare the compositional similarities.

We see that although naive methods have trade offs, our method outperforms on both metric.