Proposal: Convolutional Dictionary Learning for Arrhythmia Classification

Anthony Degleris and Eajer Toh

February 2020

1 Background

In class we discussed the *blind deconvolution* problem: given an observed image $y \in \mathbf{R}^P$ and a prior $\Gamma : \mathbf{R}^P \to \mathbf{R}_+$, find $d \in \mathbf{R}^L$ and $x \in \mathbf{R}^P$ such that $y \approx d * x$ and $\Gamma(x)$ is minimized. We quantify how well $d * x$ approximates y using a loss function $L : \mathbf{R}^P \to \mathbf{R}_+$ and using $L(y - d * x)$ as a measure of error. Blind deconvolution is useful in imaging problems where we don't fully understand the underlying generative model—like recovering an photograph taken with some unknown lens.

A natural extension of the blind deconvolution problem is to allow the observed image to come from a series of convolutions. In the *convolutional dictionary learning* problem, $y \approx \sum_{k=1}^{K} d_k * x_k$, where $d_k \in \mathbf{R}^L$ and $x_k \in \mathbf{R}^P$. Convolutional dictionary learning (CDL) allows an the observed image to have a parts based representation. Instead of being altered by a single source, the output now originates from multiple sources altered by multiple point spread functions.

We are interested in using CDL as an interpretable, unsupervised approach for learning the components of a series of images $y^{(1)}, \ldots, y^{(N)}$. For example, we could represent every frame in a video as the sum of a background component and a small number moving components that are translated across the scene over time. Another application is heartbeat analysis. We could decompose a heartbeat spectrogram into the phases of the cardiac cycle—and detect an arrhythmia when rare deviations are expressed.

1.1 Prior Work

Convolutional dictionary learning has been applied to a number of computer vision tasks. Canonical problems include classification and reconstruction. More interesting problems range from multimodal imaging [2] to super-resolution [3].

The CDL problem is often formulated with ℓ-1 regularization on each x_k, and accordingly called *convolutional sparse coding*. Efficient methods for solving CDL problems include leveraging the Alternating Direction Method of Multipliers ([1]) and exploiting frequency representations [4].

2 Proposal

We propose applying convolutional dictionary learning to arrhythmia characterization. An *arrhythmia*, defined broadly, is any ill-synchronized contraction or beating of the heart—often increasing the chances of blood pooling and clotting. These abnormal heartbeats are identifiable in electrocardiogram signals. In particular, minute changes in frequency and amplitude correspond to distinct overlaid efferent signals from multiple sources.

These frequency and amplitude differences are visually identifiable in the short-time Fourier Transform of EKG signals. The resulting time-frequency heatmap generates a 2D image which can be decomposed using convolutional dictionary learning. We propose to then apply CDL to identify and categorize the unique components of the nearly 26 different types of arrhythmia. The learned dictionary representation can then be used as the input to a classifier to quickly identify arrhythmia online. This proposed method has two benefits compared to existing techniques for arrhythmia classification. First, it is computational efficient and can be implemented in a small device (e.g. a cellphone); once the dictionary components are learned, classification only requires a single least-squares. Second, CDL provides an intuitive, visual-based pattern matching algorithm. The dictionary components are highly interpretable can be verified by a cardiologist for accuracy, unlike more abstruse classifiers.

While we hope to apply convolutional dictionary learning to arrhythmia detection, we are open to other applications of CDL. For example, we could apply CDL for video segmentation, identifying distinct spatio-temporal regions in a video.

2.1 Timeline

We'll break up our project into the following milestones:

- **Thursday, February 20th (Week 7).** Decide on a concrete objective. Identify our primary dataset.
- **Tuesday, February 25th (Week 8).** Implement preliminary versions of our algorithms. Test code on synthetic data and primary dataset. Identify limitations and weaknesses.
- **Tuesday, March 3rd (Week 9).** Address model limitations and weaknesses. Optimize results on primary dataset.
- **Wednesday, March 11th (Week 10).** Complete poster and present. Clean up code and technical notes.
- **Tuesday, March 17th (Finals Week).** Complete and submit final report. Celebrate St. Patrick's Day.

References

[1] Hilton Bristow, Anders Eriksson, and Simon Lucey. "Fast convolutional sparse coding". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2013, pp. 391–398.

[2] Kévin Degraux et al. "Online convolutional dictionary learning for multimodal imaging". In: *2017 IEEE International Conference on Image Processing (ICIP)*. IEEE. 2017, pp. 1617–1621.

[3] Shuhang Gu et al. "Convolutional sparse coding for image super-resolution". In: *Proceedings of the IEEE International Conference on Computer Vision*. 2015, pp. 1823–1831.

[4] Felix Heide, Wolfgang Heidrich, and Gordon Wetzstein. "Fast and flexible convolutional sparse coding". In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2015, pp. 5135–5143.