A Theoretical Look at the Application of Shape Analysis to Acoustics and Signals

Alex Townsend, Peter Beerli, Anke Meyer-Baese, Dennis E. Slice †

Florida State University, Department of Scientific Computing

Abstract

Over the past several years, a small community of researchers has arisen to develop and apply statistical Shape Analysis and Morphometrics to the study of Acoustics and other Time-Varying signals. Here, we discuss the theoretical underpinnings of Shape Analysis and Morphometrics as it applies to Acoustics and Signals. Ultimately, we show in this work that direct theoretical parallels can be drawn from the concepts and techniques of Morphometrics to generate meaningful and valuable results in the field of Bioacoustics. We discuss how this theoretical framework differs from previous work and provide new insights into Bioacoustics and acoustic ecology. Using this theoretical framework, we build a pipeline for analyzing bioacoustic data in a Morphometric context, and we show that the pipeline can be highly scalable to large problems. We also discuss algorithmic and implementational aspects of this pipeline including those for the extraction, alignment and analysis of individual pulses in a time-varying signal with pulse-like properties.

Introduction and Foundations

- Signals consist of a stream of data for which each data point is often correlated to its neighbors. Usually, they vary in response to some external stimulus or variable, such as the passage of time (2). Sound is one such signal.
- We assume here that all signals under consideration are of finite length and may only be sampled to an arbitrary but not infinite precision.

Definition 0.1. *A **Call** is any set of n ordered pairs $\{(k_1, t_1), ..., (k_n, t_n) | k_i, t_i \in \mathbb{R}\}$ where k_i represents an amplitude of a pressure wave in a fluid medium sampled at a specific time, t_i, which represents a time such that $t_1 < t_2 < ... < t_n$.*

Definition 0.2. *A **Pulse** is any subset, P of length m of a call C such that within a given window of size ω, noise mean μ_N and noise variance σ_N, then P is* $\{C_i, ..C_{i+m} | C_i \in C, i, m, \omega \in [1, .., ||C||], |\mu(C_i, ..C_{i+\omega})| > |\mu_N|, |\sigma(C_i, .., C_{i+\omega}| > \sigma_N))|\}$

Definition 0.3. *A **Shape** is a set of n-dimensional points $X_1, ...X_n$ ordered by an index i such that i is strictly increasing. It is the information left in these points when rotation, translation and scale are removed (5).*

Theorem 0.4. *In the case of a linear two-dimensional shape (where $x_1 < x_2 < ... < x_n$), the definitions of Pulse and Shape are logically equivalent.*

Proof. Given a two dimensional shape and a single pulse, the shape consists of a finite set of points, called landmarks, with x-position x, y-position y and index i. The values x, y may be any real numbers and i may be any positive integer. The pulse consists of a finite set of points with position along the t-axis, t, an amplitude a and an index i. The values of t, a may be any real number. In both cases, the ordering of i must be preserved for the definition to remain true. Thus if t_i is strictly increasing, t, a is simply a relabeling of the axes x, y and thus they are equivalent. Therefore, a pulse is just a 2D shape. □

- This fact lets us consider pulses to *be* shapes. Thus, the tools of Morphometrics apply.

Pulse Extraction

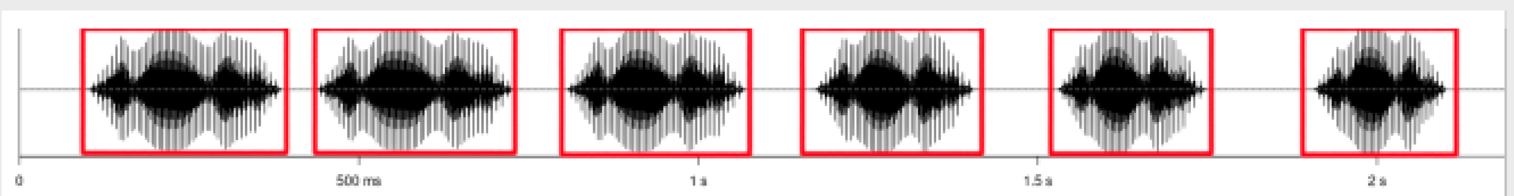

Figure 1: An example how the pipeline finds and extracts individual pulses (red boxes). The call from which these pulses were generated was simulated using the SoundGen Package in R.

Pulse Extraction Theory

- The definition of the pulse leads nicely to a method to extract pulses from a call. For each call in your set of calls, loop over every neighborhood of size ω with or without overlap, if the mean and variance exceed the tolerance, then add that to the current pulse being built. Once the mean and variance have fallen below the tolerance for enough window steps, finalize the pulse. If there is more than one pulse in a call, there will be multiple pulses in the list of pulses for each call. This will be $O(N^2)$ in the worst case.

Efficiency and Scalability

- Single-Pulse runtime is $O(N^k logN)$ where k depends on the clustering algorithm chosen.
- GPA and the Pulse Extraction algorithm are nearly embarassingly parallel.
- Both are capable of mutli-level parallelism as well (multicore and multi-node).
- Thus, several terms in the runtime analysis can be divided by (M*P) where M is the number of nodes available and P the number of processors.

Pulse Alignment

- In order to accurately measure the difference between n shapes, the shapes must be aligned such that each point in shape i is closer to it's corresponding, homologous point in shape j than it is to any other non-homologous point in shape j (5).
- Since we know that shapes and pulses are just different words for the same thing, the Generalized Procrustes Analysis method may be applied here to optimally align n pulses to one another (5). GPA is $O(N^2 logN)$ in the worst case (1).

Pulse Alignment

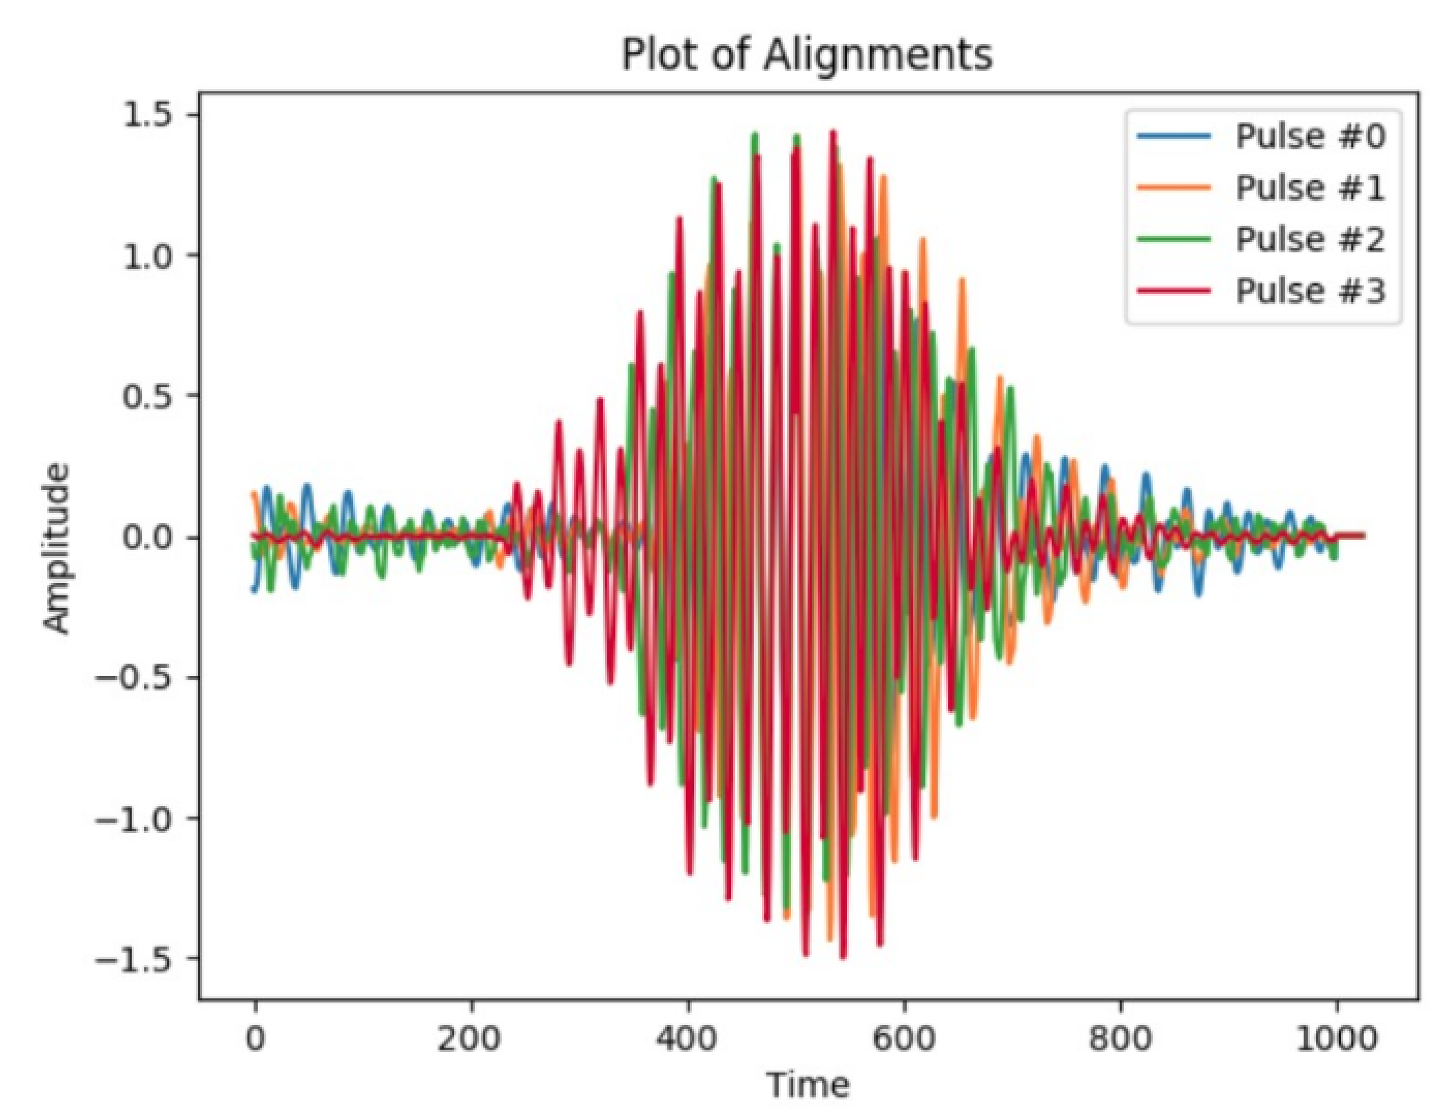

Figure 2: An example of a set of 4 maximally aligned pulses from a sample of real Frog Calls obtained from the MacAulay Library of the Cornell Lab of Ornithology.

Analysis

- Once the pulses are maximally aligned as shown in Figure 2, then we may compute the shape distance between all pairs of aligned pulses, as would be done in Morphometrics, thus generating an NxN distance matrix where N is the number of aligned Pulses to be analyzed (4).
- We may use standard techniques to analyze this distance matrix such as Neighbor-Joining Clustering, in order to generate a tree of shape distances with known branch lengths such that dialect groupings may be seen (3). Neighbor-Joining is $O(N^3)$ in the worst case (3). Most of the commonly used clustering methods are also $O(N^2 logN)$ or $O(N^3)$ (4).

Refs

1 Cormen, Thomas H., et al. Introduction to algorithms. MIT press, 2009.

2 Stiber, Michael D, Stiber, Bilin Zhang and Larson, Eric C. Signal Computing: Digital Signals in the Software Doman. 2020.

3 Studier, James A., and Karl J. Keppler. "A note on the neighbor-joining algorithm of Saitou and Nei." Molecular biology and evolution 5.6 (1988): 729-731.

4 Tan, Pang-Ning, Michael Steinbach, and Vipin Kumar. "Data mining introduction." (2006).

5 Zelditch, Miriam Leah, Donald L. Swiderski, and H. David Sheets. Geometric morphometrics for biologists: a primer. academic press, 2012.