

U. S. Naval Research Laboratory

Washington, DC 20375-5320



NRL/7160/FR--2022/1

Investigating Practical Underwater Acoustic Signal-Processing Applications of Information Geometry

DANIEL J. BROOKER

*Acoustic Signal Processing and Systems Branch
Acoustics Division*

December 16, 2022

REPORT DOCUMENTATION PAGE

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION

1. REPORT DATE December 16, 2022	2. REPORT TYPE NRL Formal Report	3. DATES COVERED	
		START DATE October 2018	END DATE October 2020

4. TITLE AND SUBTITLE
Investigating Practical Underwater Acoustic Signal-Processing Applications of Information Geometry

5a. CONTRACT NUMBER	5b. GRANT NUMBER	5c. PROGRAM ELEMENT NUMBER NISE
----------------------------	-------------------------	---

5d. PROJECT NUMBER	5e. TASK NUMBER	5f. WORK UNIT NUMBER 1L38
---------------------------	------------------------	-------------------------------------

6. AUTHOR(S)
Daniel J. Brooker

7. PERFORMING ORGANIZATION / AFFILIATION NAME(S) AND ADDRESS(ES) Naval Research Laboratory, Code 7167 4555 Overlook Ave SW Washington, DC 20375-5320	8. PERFORMING ORGANIZATION REPORT NUMBER NRL/7160/FR--2022/1
--	--

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Naval Research Laboratory, Code 7167 4555 Overlook Ave SW Washington, DC 20375-5320	10. SPONSOR / MONITOR'S ACRONYM(S) NRL-NISE	11. SPONSOR / MONITOR'S REPORT NUMBER(S)
---	---	---

12. DISTRIBUTION / AVAILABILITY STATEMENT
DISTRIBUTION STATEMENT A: Approved for public release; distribution is unlimited.

13. SUPPLEMENTAL NOTES
Karles Fellowship

14. ABSTRACT
In this work, three novel signal-processing techniques for detection, matched field processing (MFP), and depth discrimination are applied to data from the SWellEX-96 experiment. The novelty in all three cases stems from the use of information geometry to derive processors that use full-rank, cross-spectral density matrices (CSDM). Information geometry is a unique approach to probabilistic reasoning wherein distributions are viewed as points on a Riemannian manifold. This allows for hypothesis testing and for the reformulation of parameter-estimation problems using the geometric notion of distances between distributions. Five metrics are used in this work to evaluate distances between CSDMs and for the case of detection, it is found that all of the distances fail to produce a better detector than the more conventional Kullback-Leibler divergence. For MFP and depth discrimination, the Euclidean, Wasserstein-2, and root-Euclidean metrics prove successful while the Fisher-Rao and Log-Euclidean metrics fail to produce reliable results.

15. SUBJECT TERMS
Underwater acoustics, signal processing, Riemannian geometry

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT SAR	18. NUMBER OF PAGES 29
a. REPORT U/U	b. ABSTRACT U/U	c. THIS PAGE U/U		

19a. NAME OF RESPONSIBLE PERSON Daniel J. Brooker	19b. PHONE NUMBER (Include area code) 202-404-4665
---	--

This page intentionally left blank

CONTENTS

EXECUTIVE SUMMARY	E-1
1. INTRODUCTION	1
1.1 Riemannian Geometry	1
1.2 Data Processing	5
2. APPLICATION I: DETECTION	9
3. APPLICATION II: MATCHED-FIELD LOCALIZATION.....	12
4. APPLICATION III: DEPTH DISCRIMINATION	16
5. CONCLUSIONS	17
REFERENCES	18

FIGURES

1	A map showing the track of the RV Sproul during the SWelLEX-96 S5 event. The array data was taken at the point labeled “VLA.” Image courtesy http://swellex96.ucsd.edu/s5.htm , last accessed February 2021.	6
2	A spectrogram showing the signal band in the range 50 to 400 Hz for the S5 event. Vertical lines represent tones that were emitted by the two sources used in the experiment. As can be seen by the gaps in the tones, the source was turned off at the start, midway point, and near the end of the trial.	6
3	A sample of the sound-speed profile along with the bottom model used for producing replicas. Image courtesy http://swellex96.ucsd.edu/environment.htm , last accessed February 2021.	8
4	Left: The eigenspectra plotted on a log scale for the particular data-replica pair prior to adding noise. Right: Eigenspectra for the same two matrices after loading the replica with sampled noise.	9
5	ROC curves showing detector performance for the various choices of Riemannian distance. The light-gray line labeled “equality” represents probability of detection equal to probability of false alarm and is included to give a sense of the logarithm scaling used here.	11
6	ROC curves showing detector performance for the various choices of Riemannian distance and using Riemannian averages to form a null hypothesis. The light-gray line labeled “equality” represents probability of detection equal to probability of false alarm and is included to give a sense of the logarithm scaling used here.	12
7	Localization results for the five distance processors along with the Bartlett processor on the deep source and with no noise added. Solid lines on the range-versus-time curves show the ground truth as determined by GPS, while colors indicate geometric distance.	13
8	Localization results for the five distance processors along with the Bartlett processor on the shallow source and with no noise added. Solid lines on the range-versus-time curves show the ground truth as determined by GPS, while colors indicate geometric distance.	14
9	Localization results for the five distance processors along with the Bartlett processor on the deep source. Empirically estimated ambient noise has been added to the replicas to stabilize them. Solid lines on the range-versus-time curves show the ground truth as determined by GPS, while colors indicate geometric distance.	14
10	Localization results for the five distance processors along with the Bartlett processor on the shallow source. Empirically estimated ambient noise has been added to the replicas to stabilize them. Solid lines on the range-versus-time curves show the ground truth as determined by GPS, while colors indicate geometric distance.	15

TABLES

1	List of Information Distances	3
2	Source and Noise Frequencies From the SWelLEX-96 Experiment	7
3	Depth Discrimination Results Using One Cluster Each	17

This page intentionally left blank

EXECUTIVE SUMMARY

In this work, three novel underwater acoustic signal-processing techniques for detection, localization, and source classification are applied to data from the SWellEX-96 experiment. The novelty in all three cases stems from the use of information geometry, a unique, new approach to signal processing that takes advantage of the underlying geometry of acoustic data. These geometric processors are introduced in the context of underwater acoustics to make signal processing applications reliable in cluttered and low-SNR environments. Five choices of geometric distance are used in this study for all three tasks. In the case of detection, it is found that all of the distances fail to produce a better detector than the more conventional approach. For localization and source classification, three of the geometric processors prove successful while the other two fail to produce useful results.

This page intentionally left blank

INVESTIGATING PRACTICAL UNDERWATER ACOUSTIC SIGNAL-PROCESSING APPLICATIONS OF INFORMATION GEOMETRY

1. INTRODUCTION

Recently, deep connections have been revealed between the seemingly far-removed fields of differential geometry, probability, and information theory. The key insight that leads to this transformative development is the idea that families of probability distributions can be viewed as a differential manifold, a space that locally resembles Euclidean space but globally is less constrained. The advantage of this identification is that the probability manifold, in turn, can be endowed with a Riemannian metric that allows intuitive geometric notions such as distances and angles to be computed between distributions. This newly emerging subject, dubbed “information geometry,” offers the enticing opportunity for signal processing to go beyond conventional statistical techniques into what is now referred to as “geometric signal processing.” In this report, several practical applications of geometric signal processing are tested on real world data from the domain of underwater acoustics. In particular, geometrically inspired approaches are tried for the well-known problems of detection, source localization, and depth discrimination. The probability manifold in all three cases is the family of complex, multivariate Gaussian distributions, which can be used to represent data gathered on an acoustic array. Riemannian distances then are used for hypothesis testing and model evaluation in lieu of the usual statistical methods. In all three applications, the results are compelling, and yet it is unclear that there is real gain in using these more computationally intensive geometric approaches.

The report is organized as follows: The next section of this chapter will explore core concepts from Riemannian geometry including the geometric notions of a manifold, metrics, and geodesics, which will be necessary for formulating the various processors used in this study. In addition, explicit formulae will be given for the distances between points and means of data sets using Riemannian metrics. After laying out the theoretical preliminaries, the chapter will close by examining the particular data used in this study. All three of the signal-processing applications will be tested using the SWelLEX-96 data set and a section will be devoted to processing this data. Information about the environment, array processing, and modeling all will be given in this section. The next three chapters will give a description and results for each of the three applications. In the conclusion, attention will be turned to resolving the underlying issue of identifying systems for which geometric signal processing provides demonstrable benefit, and a path forward is mapped out specifically for ocean acoustics.

1.1 Riemannian Geometry

It has been known for some time that the statistical question of how to compare two models can be given an underlying *geometric* interpretation. The mathematical details of such analogies are derived rigorously in [1–5] and references therein. The distance between two probability distributions is computed by first defining a metric on the space of all probability distributions within some family and then solving the geodesic equation for that metric. If an empirically estimated probability distribution is found to be closer to some theoretical distribution in the geometric sense, then it is more likely in a probabilistic sense that

this model gave rise to the data in question. In the field of information geometry, the conventional choice of metric is the Fisher metric, which is given by minus the Hessian of the log-likelihood for a family of models. Concretely, if we have a distribution $p(x|\theta_i)$ that is characterized by a set of parameters θ_i , then the Fisher metric is

$$g_{ij} = -\mathbf{E} (\partial_i \partial_j \log (p(x|\theta_i))) . \quad (1)$$

A Riemannian metric like the one above provides a jumping-off point for generalizing all of the usual geometric notions of distance, angle, area, and volume to spaces of arbitrary curvature. This is because the metric defines a line element at the infinitesimal level according to

$$ds^2 = \sum_{ij} g_{ij} d\xi^i d\xi^j \quad (2)$$

where ds is the line element and $d\xi^i$ are a set of coordinate bases [6]. To get a clearer understanding of this, it helps to look at two-dimensional Euclidean space. The coordinate bases are dx and dy and the metric is the 2×2 identity matrix. Thus, the line element is given by

$$ds^2 = dx^2 + dy^2, \quad (3)$$

which is the Pythagorean relation in infinitesimal form. Integration allows one to move from microscopic to macroscopic. In Riemannian geometry, this can be accomplished by defining distance along a parameterized curve using the infinitesimal metric as

$$D = \int_{t_1}^{t_2} \sqrt{\sum_{ij} g_{ij} \frac{d\xi^i}{dt} \frac{d\xi^j}{dt}} dt . \quad (4)$$

Again, returning to the two-dimensional example and taking our curve as a straight line ($x(t) = t$, $y(t) = t$), the derivatives are both unity and the integral is trivial. The distance traveled over a time Δt is then $D = \sqrt{2}\Delta t$. Identifying $\Delta t = \Delta x = \Delta y$, this becomes $D = \sqrt{\Delta x^2 + \Delta y^2}$, as expected. The generalization of a straight line in Riemannian geometry is a geodesic, which is a curve that minimizes the distance integral shown above for a given metric. Geodesics on the globe, for example, are great circles, such as the equator or the prime meridian. The geodesics are determined by the metric according to the geodesic equation, a nonlinear second-order system of differential equations derived from Eq. (4) using variational techniques [6]. In principle, a distance formula like the Pythagorean formula derived above can be found for an arbitrary metric by solving these geodesic equations with fixed endpoints, but in practice, this is usually intractable, barring symmetries.

The connection between information geometry and underwater acoustics comes from recognizing that the observed cross-spectra density matrix (CSDM) of a signal on an array can be thought of as a point on one of these Riemannian manifolds. The question of comparing two signals or comparing received data with a hypothesis then can be given a geometric interpretation by using a Riemannian distance as the method of comparison [7–9]. Simply put, the more similar two signals are, the closer they will be in distance. In this case, the manifold is the space of all Hermitian positive definite matrices. Probabilistically, this is the identical to the manifold of all multivariate Gaussian distributions with fixed mean. There is a large amount

of literature about the use of different Riemannian metrics for computing distances between CSDMs, and this study will make use of four such metrics in particular, along with a Euclidean (flat) metric. The distance formulas in these cases already are known in the literature and have been well studied [5, 10–15]. The distances used in this study are given in Table 1.

Table 1—List of Information Distances

Metric	Expression
Euclidean	$\sqrt{\text{Tr} (\Sigma_1^2 + \Sigma_2^2 - 2\Sigma_1\Sigma_2)}$
Fisher-Rao	$\sqrt{\text{Tr} \log^2 \Sigma_2^{-1}\Sigma_1}$
Wasserstein-2	$\sqrt{\text{Tr} [\Sigma_1 + \Sigma_2 - 2(\Sigma_1\Sigma_2)^{1/2}]}$
Log-Euclidean	$\sqrt{\text{Tr} (\log \Sigma_1 - \log \Sigma_2)^2}$
Root-Euclidean	$\sqrt{\text{Tr} \left(\Sigma_1 + \Sigma_2 - 2\Sigma_1^{\frac{1}{2}}\Sigma_2^{\frac{1}{2}} \right)}$

While not immediately obvious, in all four cases, the Riemannian distances in Table 1 have the following three properties that make them true distance measures: First, the four distances are always greater than or equal to zero, with zero distance implying $\Sigma_1 = \Sigma_2$, second, they are all symmetric with respect to interchange of Σ_1 and Σ_2 , and third, they obey the triangle rule.

One aspect of information geometry that confuses the whole issue of selecting a metric is that there is no physical motivation behind the derivation of the various distance formulas in Table 1. In spite of this, it is still interesting to understand how the distance formulas are derived and how they have been applied elsewhere in the literature. The Fisher-Rao distance corresponds to the Fisher metric in Eq. (1) applied to the family of zero-mean Gaussian distributions. The distance formula is derived from the geodesic equation explicitly when one end point is set as the identity matrix and this can be mapped via similarity transformation into the distance between any two positive matrices (full-rank CSDMs, in this case) [12]. This distance has been used in signal-processing applications for radar detection and sonar detection, and also has been considered for matched-field processing [7, 9, 16]. The Wasserstein-2 distance is derived via a mapping between the space of positive matrices and the space of Hermitian matrices, which has a Euclidean distance measure. Both a metric and transformation are chosen such that the mapping is distance-preserving, allowing the new Riemannian distance to be derived from the Euclidean one [10, 17]. It has been applied to narrow-band sonar detection and matched-field processing by several authors [9, 11, 16, 18]. The log-Euclidean and root-Euclidean metrics both are related to the Euclidean metric, as their names imply, by a simple transformation. For the log-Euclidean distance, it is readily seen that by taking $S = \log \Sigma$, the distance (and indeed the metric, itself) becomes Euclidean [14]. The same is true of root-Euclidean with the identification $S = \Sigma^{\frac{1}{2}}$ [10, 11]. The root-Euclidean metric has been used for narrow-band sonar detection and for matched-field processing applications [9, 11], and the log-Euclidean metric has been applied to diffusion tensor imaging [14]. Finally, it is important to note that the Fisher-Rao and log-Euclidean distances are identical when Σ_1

and Σ_2 commute, and the same holds true between Wasserstein-2 and root-Euclidean. These two metrics are especially interesting from a practical standpoint because the relation to the Euclidean metric makes it easier to perform computations such as the Riemannian average, as will be shown now.

In addition to comparing data using distance formulas, there are other data analysis tasks that can be given a Riemannian interpretation. In particular, the notion of averaging has a Riemannian interpretation as well and will be used in this work both for detection and for depth discrimination. This geometric interpretation of an average comes from the fact that the arithmetic mean of a set of points,

$$\bar{x} = \frac{1}{N} \sum_{n=1}^N x_n, \quad (5)$$

can be shown to be the solution to the following optimization problem:

$$\min_x \sum_{n=1}^N (x - x_n)^2. \quad (6)$$

Identifying the quantity $(x - x_n)^2$ with the squared Euclidean distance then leads to the notion of a Frechet mean:

$$\bar{\Sigma} = \operatorname{argmin}_{\Sigma} \sum_{n=1}^N d^2(\Sigma, \Sigma_n) \quad (7)$$

where d is any Riemannian distance. For the distances listed in Table 1, analytic results for the Frechet mean of N matrices are known in only three cases. The Euclidean case has the expected result:

$$\bar{\Sigma}_E = \frac{1}{N} \sum_{n=1}^N \Sigma_n. \quad (8)$$

For the cases of log-Euclidean and root-Euclidean, the mean can be derived using the transformations mentioned previously. For the log-Euclidean distance, the Frechet mean is [14]

$$\bar{\Sigma} = \exp\left(\frac{1}{N} \sum_{n=1}^N \log \Sigma_n\right), \quad (9)$$

while for root-Euclidean, it is [11]

$$\bar{\Sigma} = \left(\frac{1}{N} \sum_{n=1}^N \Sigma_n^{\frac{1}{2}}\right)^2. \quad (10)$$

For the cases of Fisher-Rao and Wasserstein-2, the means must be computed using an iterative procedure. In both cases, an initial guess for the mean can be set by the Euclidean formula, Eq. (8). Then subsequent

iterations are found using a Newtonian gradient descent formula. For Fisher-Rao, the iterations are generated by [15]

$$\Sigma_{i+1} = \Sigma_i^{\frac{1}{2}} \exp \left(\frac{1}{N} \sum_{n=1}^N \log \left(\Sigma_i^{-\frac{1}{2}} \Sigma_n \Sigma_i^{-\frac{1}{2}} \right) \right) \Sigma_i^{\frac{1}{2}}, \quad (11)$$

where the Σ_n are the data matrices and the Σ_i are the estimates of the average. For Wasserstein-2, the iterations are generated by [13]

$$\Sigma_{i+1} = \Sigma_i^{-\frac{1}{2}} \exp \left(\frac{1}{N} \sum_{n=1}^N \left(\Sigma_i^{\frac{1}{2}} \Sigma_n \Sigma_i^{\frac{1}{2}} \right)^2 \right) \Sigma_i^{-\frac{1}{2}} \quad (12)$$

In both cases, it has been proven that this algorithm in the limit of an infinite number of steps converges to the Frechet mean defined by Eq. (7) [13, 15]. Now that the mathematical preliminaries have been established, the next section will examine a particular data set, the SWellEX-96 experiments S5 event [19, 20], which will be explored using Riemannian signal-processing techniques.

1.2 Data Processing

The SWellEX-96 S5 data set [19, 20] was chosen for this study because it has been used previously for all the signal-processing tasks under consideration: detection, localization, and depth discrimination [21–23]. The experiment was conducted in May of 1996 several kilometers offshore from San Diego. During the S5 event, two sources were towed simultaneously by the same ship at a speed of about 5 knots for a total duration of 75 minutes. The deep source was towed at a depth of approximately 54 m and the shallow source at a depth of about 9 m. The ship track and the location of the acoustic arrays are shown in Fig. 1. All of the trials in this study will be done using data collected on the vertical array labeled “VLA” in Fig. 1. The array was bottom-moored at a depth of 216.5 m and had 21 hydrophones with 6 m spacing between them. Phone data was sampled at a rate of 1.5 kHz and a fast Fourier transform (FFT) of length 4,096 was used to extract complex pressure values with a snapshot length of approximately 2.73 s and a frequency bin size of 0.3662 Hz. The snapshots were buffered together so that the first half of each snapshot was the last half of the previous snapshot, effectively doubling the total number. In total, there were 3,294 snapshots taken. A spectrogram showing the incoherent average power on the array during the whole experiment is shown in Fig. 2.

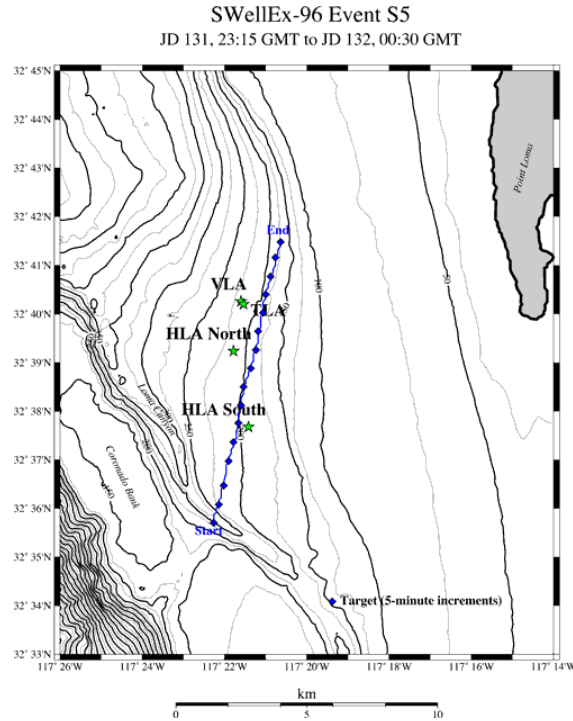


Fig. 1—A map showing the track of the RV Sproul during the SWellEX-96 S5 event. The array data was taken at the point labeled “VLA.” Image courtesy <http://swellex96.ucsd.edu/s5.htm>, last accessed February 2021.




Fig. 2—A spectrogram showing the signal band in the range 50 to 400 Hz for the S5 event. Vertical lines represent tones that were emitted by the two sources used in the experiment. As can be seen by the gaps in the tones, the source was turned off at the start, midway point, and near the end of the trial.

Of the tonal set that was broadcast, several can be observed clearly in Fig. 2 as vertical lines in the spectrogram. Of the broadcast tones, the tonals to be processed for localization and depth discrimination along with several noise frequencies are displayed in Table 2.

Table 2—Source and Noise Frequencies From the SWelLEX-96 Experiment

Source	Frequencies (Hz)
Deep source	49, 64, 79, 94, 112, 130, 148, 166, 201, 235, 283, 338, 388
Shallow source	109, 127, 145, 163, 198, 232, 280, 335, 385
Noise	92, 107, 125, 143, 161, 179, 248, 296, 351, 398

From the complex spectra, CSDMs were generated as follows: The full data set was split up into segments of 28 snapshots, with each segment overlapping half of the previous segment. This led to a total of 232 segments of data. Then for each segment, the CSDM at each frequency is computed as

$$\Sigma(f)_{ij} = \frac{1}{N-1} \sum_{t=1}^{28} \tilde{p}_{ti}^* \tilde{p}_{tj} . \quad (13)$$

The localization and depth-discrimination processors are both generalizations of matched-field processing and require replica CSDMs [9]. Replica CSDMs were made on a 216-m-by-10-km grid using a spacing of 2 m in depth and 50 m in range. In order to make the replicas full rank, 100 random samples of the sound speed profile (SSP) were drawn from a normal distribution whose mean and covariance were determined empirically using an ensemble of 39 conductivity-temperature-depth (CTD) casts that were taken during the SWelLEX experiment. The acoustic field then was evaluated using the range-dependent acoustic model (RAM) [24] for each frequency in Table 2. The replica CSDM was determined from the RAM models using Eq. (13). These models all were range-independent with a water column depth of 216.5 m and a bottom profile that consisted of a 23.5-m sediment layer followed by an 800-m sub-basement and a fluid half-space basement as specified by the SWelLEX environmental database [20, 25]. A typical sound speed along with the bottom model is shown in Fig. 3.

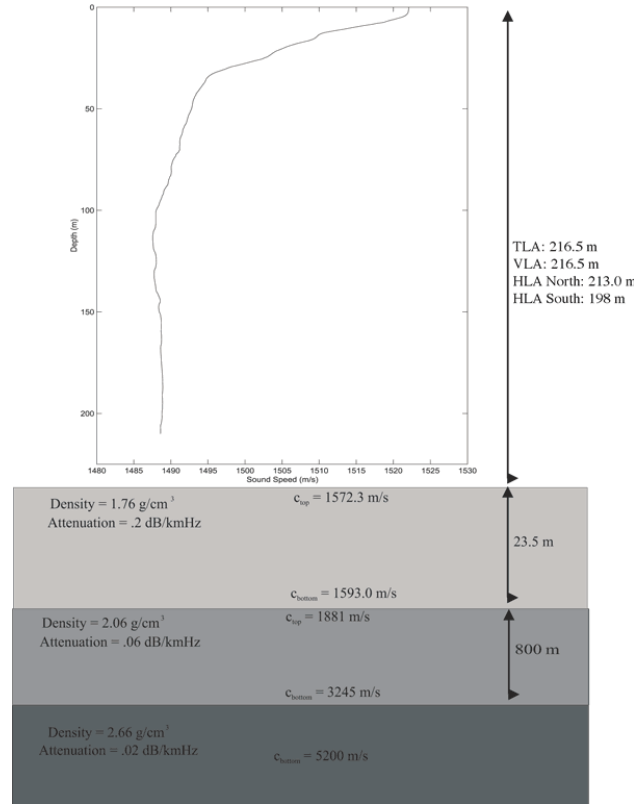


Fig. 3—A sample of the sound-speed profile along with the bottom model used for producing replicas. Image courtesy <http://swellex96.ucsd.edu/environment.htm>, last accessed February 2021.

Since the Riemannian distances and Frechet means will require matrix inversions and matrix logarithms, it is important that the replica CSDMs be numerically stable enough to support such calculations. To ensure that the replica matrices are not too singular and to improve the quality of the simulation, another replica set was made by loading the original replicas using an empirical estimate of the ambient noise CSDM. This was done by adding a scaled-data CSDM to the replica using one of the frequencies which does not contain any signal. The scaling was set at each time segment and for each frequency using an empirical estimate of the SNR based on the power spectra:

$$SNR(f_r, f_n) = 10^{(P(f_r) - P(f_n))/10}, \quad (14)$$

where P is the power in decibels, f_r is the replica frequency, and f_n is the noise frequency. Both CSDMs first had their traces normalized to unity, then the replica CSDM was scaled by

$$\Sigma_{r,scaled} = \Sigma_r (SNR - 1) / SNR, \quad (15)$$

and the noise CSDM was scaled by

$$\Sigma_{n,scaled} = \Sigma_n / SNR, \quad (16)$$

so that the new replica CSDM is

$$\Sigma_{r, new} = \Sigma_{r, scaled} + \Sigma_{n, scaled} . \quad (17)$$

For this study, the replica frequencies are shown in Table 2 along with the noise frequencies that were chosen to avoid overlapping with any of the other lines [20]. To see why this loading is important for increasing the fidelity of the replicas, an example of how the eigenspectra of a particular replica changes with this noise loading is shown in Fig. 4. As can be seen, the eigenvalues are in much better agreement, especially those at lower power. Now that both the theoretical preliminaries and the treatment of data are established, the first of three signal-processing applications, detection, will be examined.




Fig. 4—Left: The eigenspectra plotted on a log scale for the particular data-replica pair prior to adding noise. Right: Eigenspectra for the same two matrices after loading the replica with sampled noise.

2. APPLICATION I: DETECTION

Geometric detection has been explored previously in a number of works [7, 8, 11, 16]. This section in particular is based upon and uses the same methodology as Brooker et. al [16]. There are two key differences between the detectors implemented here and those typically found in textbooks on signal processing [26]. First, the detectors in this study do not perform matched filtering. Instead, they use a signal-discovery technique motivated by the relative entropy detectors of Mignerey et. al. [27]. The distinction is that instead of testing two hypotheses (H_0, H_1) corresponding to a noise distribution H_0 and a candidate waveform H_1 *a la* matched filtering, these detectors test only against a single null hypothesis. Second, the choice of test statistic in this study is based on information geometry and will use the Riemannian distances shown in Table 1. As a point of comparison, the Kullback-Leibler (KL) divergence also will be used as a more conventional choice of dissimilarity measure between distributions [28]. The KL divergence is defined for two probability

distributions (P, Q) as

$$D(P||Q) = \sum_{x_i} P(x_i) \log \frac{P(x_i)}{Q(x_i)}. \quad (18)$$

This measure is conventional because it satisfies the Neyman-Pearson lemma, thus providing a uniformly most powerful test in the limit of large sample sizes [26, 29]. For the case of two multivariate complex Gaussian distributions with zero mean (labeled 1 and 2), the KL divergence is given by [28]

$$D(\mathcal{N}_1||\mathcal{N}_2) = \frac{1}{2} \left[\text{Tr} \left(\Sigma_2^{-1} \Sigma_1 - I \right) + \log \left(\frac{\det \Sigma_2}{\det \Sigma_1} \right) \right]. \quad (19)$$

To implement this expression (or any of the distances listed in Table 1) in a detector, one assigns a candidate distribution, which may or may not contain signal, to \mathcal{N}_1 and a noise distribution to \mathcal{N}_2 that is assumed to contain no signal. If the dissimilarity (or distance) measure D between the two distributions exceeds a chosen threshold, then a detection is made.

All of the distances given in Table 1 have been implemented as detectors previously [7, 8, 11, 16]; however, such works are usually inconclusive as to which distance, if any, works best. Previously, Riemannian distances were introduced in the context of detection statistics to provide improved results in noisy, cluttered environments where signals are well modeled as zero-mean Gaussian processes [7]. However, it is difficult to judge metrics against one another since there is no *a priori* criteria for evaluating them as detectors independent of the environment and the measurement model under consideration. In addition, detectors may perform differently depending on the environment, even when used for the same application [27]. Wong et. al. have suggested that a joint Riemannian machine-learning approach may provide a resolution to this issue since a metric (i.e. detector) can be learned that maximizes the level of dissimilarity between a noise ensemble and incoming signals [11]. In their own paper, however, the authors test several versions of this machine learning approach using several different metrics as starting points. It is not clear whether the learned metrics are the same or even related in among the different trials, and still the question remains of which metric to use as a starting point. Given the confusion, the trend in the literature has been to test many different metrics for a given application. One trend that does stand out is that in studies that do make use of the KL divergence as a point of reference, the KL divergence typically performs quite well, leaving one to wonder whether detection is even a suitable problem for this Riemannian approach [8, 16].

The Riemannian detectors are implemented as follows: Two bands are extracted — one ranging from 100 to 400 Hz and one ranging from 450 to 750 Hz. As seen from the spectrogram in Fig. 2, the 100-to-400-Hz band contains a number of narrow-band signal lines, while the 450-to-750-Hz band was empty. For each time segment, the frequency bins in each band are sorted by total incoherent power and all of the bins with below-median power are used to form a noise ensemble for each band. Then a single-noise CSDM is found in each band by averaging the statistics of these ensembles, and this CSDM serves as the null hypothesis for the detector. Then the detector output is formed as a spectrum of Riemannian distances between the signal and noise CSDMs.

Performance of the detectors is evaluated by generating empirical receiver operating characteristic (ROC) curves from mutual cumulative distribution functions (CDF) using the algorithm in Mignerey et al. [27]. The CDFs required to make the ROC curves are generated from information values obtained for each time segment by summing the Riemannian distance spectrum over the whole frequency band. These information

values, one for each time segment, are used to form ensembles required by the ROC curves. Within each ensemble, data values are sorted and normalized to form a staircase function for the CDF, which is then smoothed. The empirical ROC curves are shown for each metric in Fig. 5 and the KL-divergence performs the best by a large margin.




Fig. 5—ROC curves showing detector performance for the various choices of Riemannian distance. The light-gray line labeled “equality” represents probability of detection equal to probability of false alarm and is included to give a sense of the logarithm scaling used here.

The poor performance of the Riemannian distances is unsurprising, given that similar results were found in Brooker et. al. [16]. That the KL divergence performs so much better than the other distances could be due in part to the fact that this is a rather high SNR environment, whereas Riemannian distances for detection were first introduced for cluttered, noisy environments [7]. It remains the case, however, that in any study that includes the KL divergence alongside a Riemannian distance for detection the KL divergence performs at least as well as the Riemannian distances [8, 16, 30]. One possible route to improving the geometric detectors is to use a Riemannian mean as discussed previously when averaging together the noise ensemble to form a null hypothesis as suggested by Hua et. al. [30]. This has been done and the results are shown in Fig. 6. The improvements offered are significant, but it still does not account for the performance gap between the Riemannian detectors and the KL divergence. It might be concluded, then, that detection is not a suitable application of geometric signal processing in the sense that there is little evidence for significant value added over conventional detection schemes. Our next application will be localization of the narrow band lines seen in Fig. 2 using matched-field processing.

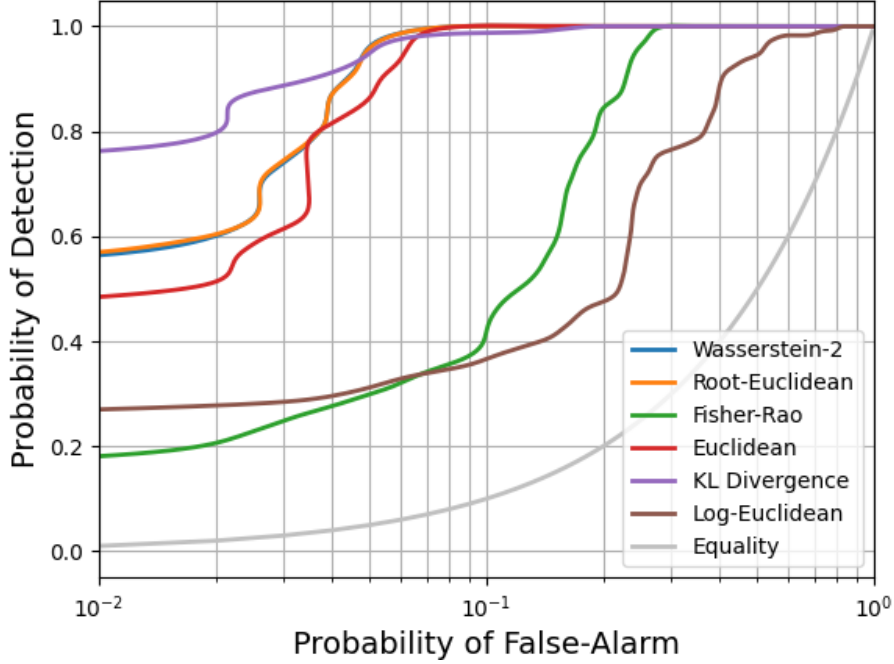


Fig. 6—ROC curves showing detector performance for the various choices of Riemannian distance and using Riemannian averages to form a null hypothesis. The light-gray line labeled “equality” represents probability of detection equal to probability of false alarm and is included to give a sense of the logarithm scaling used here.

3. APPLICATION II: MATCHED-FIELD LOCALIZATION

Source localization using an information geometric version of matched-field processing (MFP) first was proposed by Finette and Mignerey [9] and later was demonstrated in a deep-water environment by Cao et. al. [18]. The idea behind this generalization stems from the fact that the Bartlett processor, itself, has a geometric interpretation. The Bartlett processor, a starting point for many localization algorithms, takes a series of replica vectors $r(x)$ that represent a modeled array response for a source at a location x and compares it with a received data d using an inner product:

$$P_{Bartlett}(x) = \mathbf{E}(|r^H(x)d|^2) \equiv r^H(x)\Sigma r(x), \quad (20)$$

where Σ is the data CSDM [22, 31, 32]. Localization then is performed by computing the Bartlett response surface for a given CSDM and the location corresponding to the replica with the greatest response is selected. The geometric interpretation comes from recognizing that in any Euclidean space, the inner product is intimately related to geometric distance (D) by the law of cosines:

$$D(A, B)^2 = \|A\|^2 + \|B\|^2 - 2A \cdot B, \quad (21)$$

where A and B are any two points. Since MFP replicas and data are typically normalized to unity first, the Bartlett processor then could be written equivalently as

$$P_{\text{Bartlett}}(x) = \mathbf{E}(|1 - D^2(r, d)|^2). \quad (22)$$

What Finette and Mignerey did was to promote r from a replica vector to a replica CSDM, doing away with the expectation value and the squaring of the output. Then they replace the factor of $1 - D^2$ with D using a Riemannian distance and, instead of picking the replica that maximizes the inner product, they pick the replica that minimizes distance [9]. The generalization from replica vectors to replica CSDMs is quite powerful from a modeling point of view, since more physics can be included in generating replica CSDMs than replica vectors [9]. Moreover, from a signal-processing point of view, there are more degrees of freedom to be constrained when using CSDMs if adaptive constraints are to be added later [33]. In the paper by Finette and Mignerey, the additional physics came via a shallow-water linear-internal-wave model. The internal waves generated a stochastic perturbation to the sound speed that contributes to the CSDM but averages to zero in the mean [9]. The sound-speed distribution that was generated from the SWelEX experiments' CTD casts and the added noise estimate serve a similar role here.

All of the frequencies listed in Table 2 are processed individually and then the distances are averaged together to form a single incoherent output for each of the two sources. Then a location is selected by picking the replica that minimizes the incoherent output. The results are shown in Figs. 7 and 8 for the replicas with no noise added and in Figs. 9 and 10 for the replicas with empirical noise estimates. Both trials were

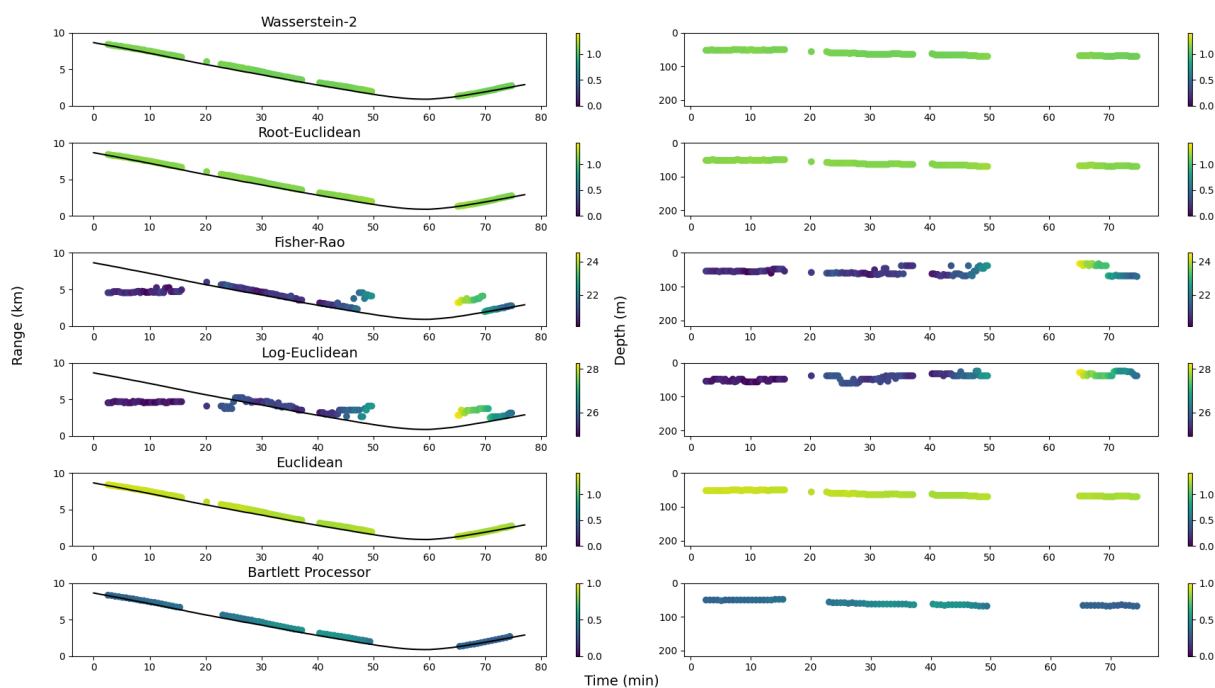


Fig. 7—Localization results for the five distance processors along with the Bartlett processor on the deep source and with no noise added. Solid lines on the range-versus-time curves show the ground truth as determined by GPS, while colors indicate geometric distance.

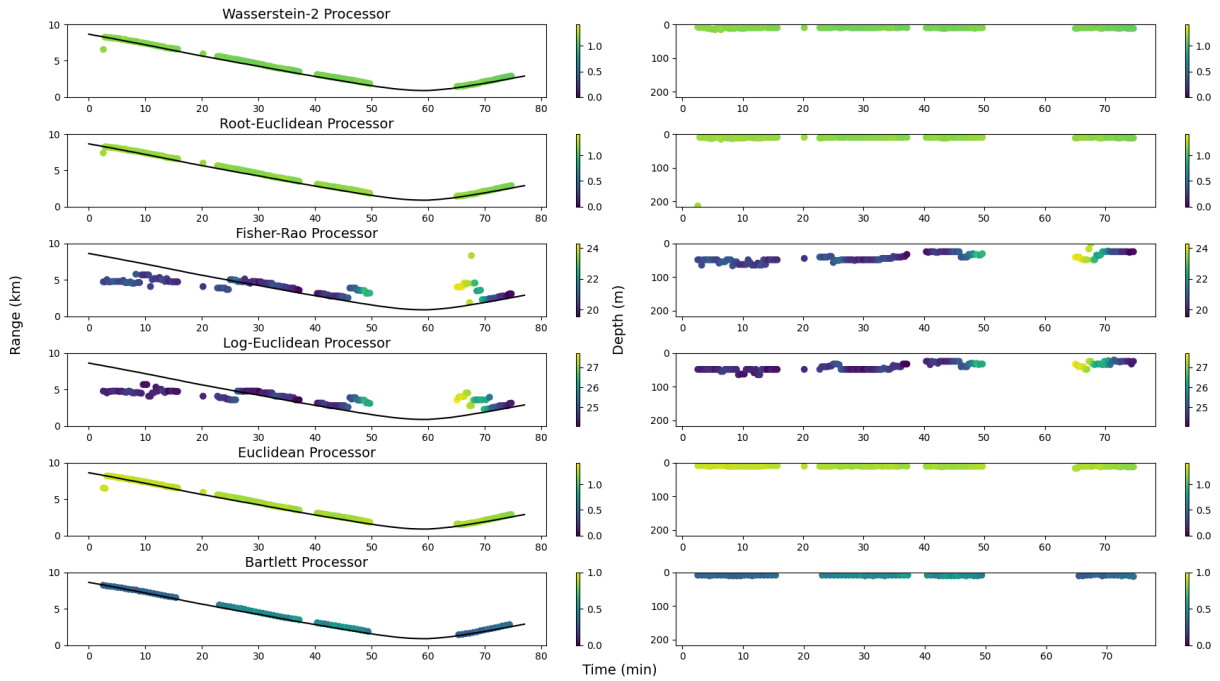


Fig. 8—Localization results for the five distance processors along with the Bartlett processor on the shallow source and with no noise added. Solid lines on the range-versus-time curves show the ground truth as determined by GPS, while colors indicate geometric distance.

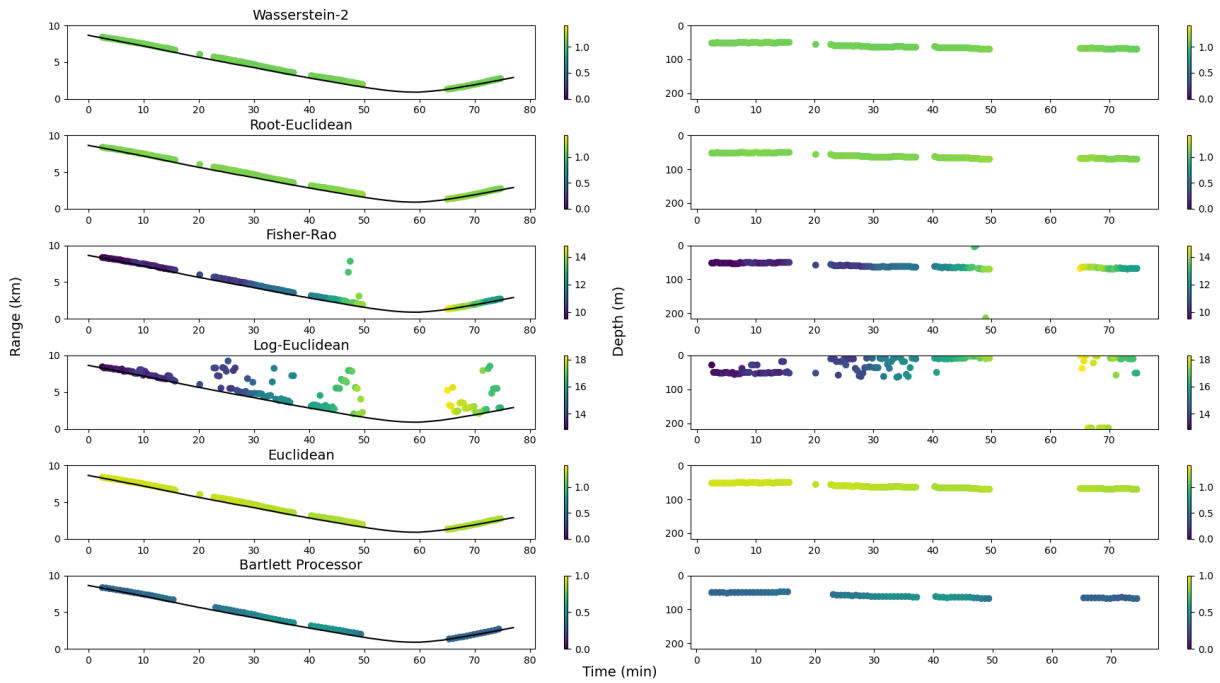


Fig. 9—Localization results for the five distance processors along with the Bartlett processor on the deep source. Empirically estimated ambient noise has been added to the replicas to stabilize them. Solid lines on the range-versus-time curves show the ground truth as determined by GPS, while colors indicate geometric distance.

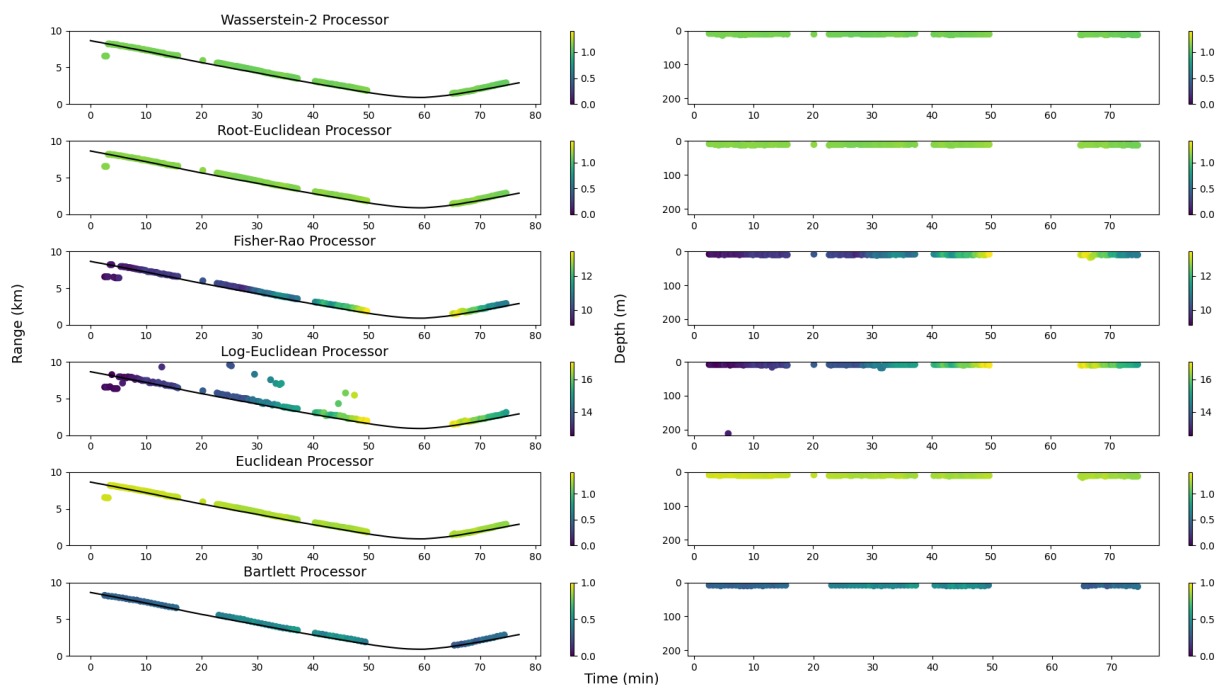


Fig. 10—Localization results for the five distance processors along with the Bartlett processor on the shallow source. Empirically estimated ambient noise has been added to the replicas to stabilize them. Solid lines on the range-versus-time curves show the ground truth as determined by GPS, while colors indicate geometric distance.

performed using the full replica grid, which has a spacing of 2 m in depth by 50 m in range and covered a 216-m-by-10-km water column. The Wasserstein-2, root-Euclidean, and Euclidean processors have localized the two sources correctly in all cases, while the log-Euclidean processor and the Fisher-Rao metric both fail when no noise is added. The failure of the Fisher-Rao metric also was noted by Cao et al., who did not give results for that processor [18].

It is important to note that the numerical output of the processor represents how similar the data and replica are, but it is unclear if there is any physical interpretation as to why it takes on the particular numerical value that it does [9]. For all of the Riemannian processors, a distance of zero would represent a perfect match. The Euclidean, Wasserstein-2, and root-Euclidean distances all have an upper bound on the distance of $\sqrt{2}$ due to the fact that the replica and data CSDMs have their traces normalized to unity. The Fisher-Rao and log-Euclidean processors have no such bound, but it would seem based on the values for the other three distances that in spite of correctly localizing the source in some cases, there is not very much agreement between the sources and the replicas. This might indicate that more modeling is needed of the environment, and that the noise loading procedure could be improved. There are a number of ways this might be done in future works. One possibility would be to include spatial variations in the sound-speed profile as was done in the original work of Finette and Mignerey [9]. With enough spatial sampling of the sound speed, this could be modeled using a Karhounen-Loeve expansion. Another option would be to use uncertainty to generate a CSDM. In this approach, replicas could be generated over a range of unknown environmental parameters and these could be averaged together. A final possibility for improving the Riemannian implementation of MFP would be to include constraints. The minimum variance constraint C_{MVDR} [34, 35], like the Bartlett

processor, itself, has a geometric interpretation in terms of an inner product,

$$C_{MVDR} \equiv w^H r - 1 = 0, \quad (23)$$

and this could be generalized using inner products defined by a Riemannian metric. This would be challenging, however, due to the fact that in Euclidean space, there is a natural correspondence between points and vectors, namely that each point can be identified with the vector from the origin to that point, whereas Riemannian manifolds have no such correspondence. Since the view has been taken that the data and replicas are to be represented by a point on the manifold, there would be some ambiguity as to which vectors would be used in forming constraints. Nevertheless, once a choice is made to resolve this ambiguity, the MFP process could be formulated as a manifold optimization problem [36]. For MVDR, the problem is to minimize the inner product between the data and the weights subject to the MVDR constraint. For Riemannian adaptive MFP, the problem would be to maximize the distance between data and weights subject to geometric constraints. This is all beyond the scope of the current work, however, and is left as a topic for future study. Attention now is turned to the last of the three signal-processing applications, source depth discrimination.

4. APPLICATION III: DEPTH DISCRIMINATION

It has been observed that in an environment with a thermocline, the mode structure of the acoustic field is such that the surface sources tend to excite only higher-order modes, since the lower-order modes are effectively trapped below the thermocline. This leads to the need for several processors for depth discrimination, with one popular choice being the so-called “trapped energy ratio” [23, 37, 38]. As a binary hypothesis test, the trapped energy ratio involves picking a certain number of modes that are trapped and observing what fraction of the total modal energy they take up for a given source. Thus, a low trapped energy ratio should correspond to a surface source and a high trapped energy ratio will indicate a submerged source. This is a fantastic idea because it takes advantage of a specific environmental signature and exploits it to perform source identification. One key drawback to these techniques, however, is that they do not work in environments with a large degree of mode mixing or in environments that lack a thermocline. An alternative to this approach that could work in any environment would be to perform a crude MFP that matches the data to replicas in depth only. If the goal is simply to determine whether a received signal is near the surface or deeply submerged, then it is possible that only a few replicas would be needed. This is the approach taken in this work and to facilitate a reduction in size of the replica set, clustering will be employed. Clustering of MFP replicas has been performed previously as a way to facilitate direction of arrival estimation, but it has not been tried yet for depth discrimination [39]. The advantage is that clustering of replicas can be performed as part of the preparation for an experiment, and then at sea, the evaluation will be much faster, since received signals are only matched against one representative from each cluster. In this case, the clusters correspond to the two source depths, but in principle, there might be more clusters if there were more depth bins being considered or if there were many types of sources being searched for (e.g. mechanical vs biologic). Clustering, in general, is a procedure that is done in two parts. First, the position of new data relative to the existing data set is evaluated. Then it is put into the cluster to which it belongs, according to some rule. The first question is a question of geometry, and so it is here that information geometry can be exploited.

The clustering method used here is a greedy algorithm based on previous work for clustering of beams in horizontal beamforming [39]. The replica grid was reduced from 2-m depth increments and 50-m range increments to 10 m by 250 m at this stage, and the replicas without any added noise were used to make

the problem more computationally tractable. Replicas with less than 20 m depth were put into the surface category, while those below 20 m depth were put into the submerged category. A matrix is computed for each metric containing the distances between all pairs of replicas. A threshold distance d_t is set, the replica with the largest number of neighbors at a distance less than d_t is chosen as the first cluster center, and all the replicas of distance $d < d_t$ from this center are assigned to that cluster. This is repeated until there is at least one cluster center containing each replica. Then another pass through the replica set is made and replicas are re-assigned to whichever cluster center is closest. Finally, a new center is formed for each cluster by taking a Riemannian mean of all the replicas in the cluster. After this training, evaluation is performed on the SWellEX data in a similar fashion. Data CSDMs are assigned to the cluster centers closest to them, and the processor output reports whether the assigned cluster is surface or submerged. For this trial, the threshold distance is set high enough that there is only one cluster each for surface and submerged. All 19 frequencies listed in Table 2 are treated individually and the percentages of correct localizations are shown in Table 3 for the Wasserstein-2, root-Euclidean, and Euclidean metrics.

Table 3—Depth Discrimination Results Using One Cluster Each

Metric	Percent correct
Euclidean	88.8
Wasserstein-2	89.8
Root-Euclidean	92.2

The log-Euclidean metrics and Fisher-Rao metrics are not reported in this case because they failed to pick the right cluster correctly with any reliability, each having a success rate of around 50 percent. This was to be expected, however, given the MFP results reported in the previous chapter. It is possible that the Fisher-Rao and log-Euclidean metrics would succeed at depth discrimination using the replica set that has added ambient noise, but these replicas feature a different noise CSDM for each time segment, meaning the clustering would have to be done each time a new data CSDM is tested. This would be far too impractical to implement and completely defeats the purpose of clustering, namely that it could be done in advance of any trials. It is quite impressive, however, that even with only two replicas, the depth discrimination is a success for the other three metrics. This should have been expected, perhaps, since depth discrimination already has been performed successfully on this data set, but the approach here is still novel and could be applied to more challenging environments, such as environments lacking a thermocline [23, 37, 38]. One issue with this whole procedure is that it is not at all clear how the distance threshold used in the clustering determines the number of clusters. This could be determined empirically, but it will be different for different replica sets. In this case, it is not very enlightening to explore, since with only one cluster each, the surface depth discrimination already was a success, but it would be interesting to explore in other data sets.

5. CONCLUSIONS

In this work, signal-processing techniques based on information geometry have been applied to three different tasks: detection, localization, and source depth discrimination. The Riemannian geometry based techniques achieved varying levels of success with the Euclidean, Wasserstein-2, and root-Euclidean distances performing the best across all three tasks. It remains to be seen, however, whether there are any tasks for which the non-Euclidean distances provide a decided advantage, given that the original motivations for geometric signal-processing applications were for low-SNR, cluttered, and highly variable environments [7–9]. The

SWelLEX environment, by contrast, is fairly stationary, and high SNR was seen in the spectrogram in Fig. 2. This makes it a great testing ground for new signal-processing techniques, but now that the viability of these Riemannian techniques has been demonstrated, they should be applied to other sites. One complication, though, is that it is hard to know *a priori* whether a given environment will need geometric signal processing. One possibility would be to use curvature as guide to understanding the need for a non-Euclidean metric. Ricci curvature, in particular, is a local measure of curvature that is specific to each metric and that varies over the manifold. On all manifolds, distances are approximately Euclidean for nearby points, but in regions of higher curvature, the Riemannian distance diverges from its Euclidean approximation more quickly [6]. By connecting curvature directly to environmental models, a study could be done to determine which ocean environments are amenable to geometric signal processing.

In addition, there are still more questions that can be answered for the particular applications studied in this work. How can the replica CSDMs be improved? Is there a better way to estimate ambient noise besides the empirical estimates used in the MFP chapter? Can the performance of the Fisher-Rao and log-Euclidean processors be improved for MFP and depth discrimination with better replicas? How should the clustering approach to depth discrimination be implemented in an environment where the naive choice of distance threshold used here is no longer sufficient? If a study of curvature in ocean environments were performed, these questions all could be explored in a situation in which information geometry provides a clearer value added.

REFERENCES

1. C. Rao, "On the distance between two populations," *Sankhya* **9**, 246 – 248 (1948).
2. J. Burbea and C. Rao, "Differential metrics in probability spaces," *Prob. Math. Statist.* pp. 115 – 132 (1982).
3. S. Amari, "Differential geometry of curved exponential families- curvature and information loss," *Ann. Stat.* **10**, 357 – 385 (1982).
4. S. Amari, "A foundation of information geometry," *Electronics and Communications in Japan* **66-A**, 1–10 (1983).
5. L. Skovgaard, "A Riemannian geometry of the multivariate normal distribution," *Scandinavian J Stats* **11**, 211 – 223 (1984).
6. C. W. Misner, K. S. Thorne, and J. A. Wheeler, *Gravitation* (San Francisco: W.H. Freeman and Co., 1973, 1973).
7. M. Arnaudon, F. Barbaresco, and L. Yang, "Riemannian Medians and Means With Applications to Radar Signal Processing," *IEEE Journal of Selected Topics in Signal Processing* **7**(4), 595–604 (2013).
8. Y. Cheng, X. Hua, H. Wang, Y. Qin, and X. Li, "The Geometry of Signal Detection with Applications to Radar Signal Processing," *Entropy* **18**(11), 381 (2016).
9. S. Finette and P. Mignerey, "Stochastic matched-field localization of an acoustic source based on principles of Riemannian geometry," *The Journal of the Acoustical Society of America* **143**(6), 3628–3638 (2018).

10. Y. Li and K. Wong, "Riemannian distances for signal classification by power spectral density," *IEEE J. Sel. Topics Signal Process.* **7**, 655 – 669 (2013).
11. K. M. Wong, J. Zhang, J. Liang, and H. Jiang, "Mean and Median of PSD Matrices on a Riemannian Manifold: Application to Detection of Narrow-Band Sonar Signals," *IEEE Transactions on Signal Processing* **65**(24), 6536–6550 (2017).
12. R. Bhatia, *Positive Definite Matrices*, Princeton Series in Applied Mathematics (Princeton University Press, New Jersey, 2007).
13. R. Bhatia, T. Jain, and Y. Lim, "On the Bures-Wasserstein distance between positive definite matrices," 2017.
14. V. Arsigny, P. Fillard, X. Pennec, and N. Ayache, "Log-Euclidean metrics for fast and simple calculus on diffusion tensors," *Magn Reson Med* **56**, 411–21 (2006).
15. X. Pennec, P. Fillard, and N. Ayache, "A Riemannian framework for tensor computing," *I. J. Computer Vision* **66**, 41 – 66 (2006).
16. D. Brooker, S. Finette, and P. Mignerey, "Detection on a Distributed Network Using Information Geometry," *The Journal of the Acoustical Society of America* (2020).
17. A. Uhlmann, "Parallel Transport and 'Quantum Holonomy'," *Rep. Math. Phys.* **24**, 229–240 (1986).
18. R. Cao, K. Yang, Y. Ma, Y. Qiulong, and Y. Shi, "Passive broadband source localization based on a Riemannian distance with a short vertical array in the deep ocean," *The Journal of the Acoustical Society of America* **145**(6), EL567–EL573 (2019).
19. N.Booth, "SWellEx-96 preliminary data Report," *NRaD/NCCOSC Code 881* (San Diego, CA, May 1996).
20. N.Booth, O. Newell, W. Hodgkiss, and D. Ensberg, "SWellEx-96 Experiment Acoustic Data," 2015. URL <http://swelllex96.ucsd.edu/environment.htm>.
21. L. Sha and L. W. Nolte, "Bayesian Sonar Detection Performance Prediction With Source Position Uncertainty Using SWellEx-96 Vertical Array Data," *IEEE Journal of Oceanic Engineering* **31**(2), 345–355 (2006).
22. N. O. Booth, P. A. Baxley, J. A. Rice, P. W. Schey, W. S. Hodgkiss, G. L. D'Spain, and J. J. Murray, "Source localization with broad-band matched-field processing in shallow water," *IEEE Journal of Oceanic Engineering* **21**(4), 402–412 (1996), doi:10.1109/48.544051.
23. V.E. Premus, J. Ward, and C. D. Richmond, "Mode filtering approaches to acoustic source depth discrimination," Proceedings of the Conference Record of the Thirty-Eighth Asilomar Conference on Signals, Systems and Computers, 2004., volume 2, 2004, pp. 1415–1420 Vol.2.
24. M. Collins, "Users Guide for Ram," oalib.hlsresearch.com/PE (Last Accessed 2019).
25. R. Bachman, P. Schey, N. Booth, and F. Ryan, "Geoacoustic databases for matched-field processing: Preliminary results in shallow water off San Diego, California," *The Journal of the Acoustical Society of America* **99**(4), 2077–2085 (1996).

26. L. Scharf and C. Demeure, *Statistical Signal Processing: Detection, Estimation, and Time Series Analysis* (Addison-Wesley, Massachusetts, 1991).
27. P. Mignerey, A. Turgut, J. Schindall, and D. Goldstein, “Evaluation of Relative Entropy for Distributed Passive Detection of Weak Acoustic Signals,” *IEEE Journal of Oceanic Engineering* **42**(1), 219–230 (2017).
28. S. Kullback, *Information Theory and Statistics* (Dover, Mineola, 1968).
29. J. Neyman and E. Pearson, “On the problem of the most efficient tests of statistical hypotheses,” *Phil. Trans. Roy. Soc. Lond. A* pp. 289–337 (1933).
30. X. Hua, Y. Chen, H. Wang, Y. Qin, and Y. Li, “Geometric Means and Medians with Applications to Target Detection,” *IET Signal Processing* **11**(6), 711–720 (2017).
31. A. B. Baggeroer, W. A. Kuperman, and H. Schmidt, “Matched field processing: Source localization in correlated noise as an optimum parameter estimation problem,” *The Journal of the Acoustical Society of America* **83**(2), 571–587 (1988).
32. A. B. Baggeroer, W. A. Kuperman, and P. N. Mikhalevsky, “An overview of matched field methods in ocean acoustics,” *IEEE Journal of Oceanic Engineering* **18**(4), 401–424 (1993), doi:10.1109/48.262292.
33. S. Czenszak and J. Krolik, “Robust wideband matched-field processing with a short vertical array,” *The Journal of the Acoustical Society of America* **101**(2), 749–759 (1997).
34. J. Capon, “High-resolution frequency-wavenumber spectrum analysis,” *Proceedings of the IEEE* **57**(8), 1408–1418 (1969), doi:10.1109/PROC.1969.7278.
35. H. Cox, R. Zeskind, and M. Owen, “Robust adaptive beamforming,” *IEEE Transactions on Acoustics, Speech, and Signal Processing* **35**(10), 1365–1376 (1987).
36. N. Boumal, B. Mishra, P. Absil, and R. Sepulchre, “Manopt, a Matlab Toolbox for Optimization on Manifolds,” *Journal of Machine Learning Research* **15**(42), 1455–1459 (2014). URL <https://www.manopt.org>.
37. E. Conan, J. Bonnel, T. Chonavel, and B. Nicolas, “Source depth discrimination with a vertical line array,” *The Journal of the Acoustical Society of America* **140**(5), EL434–EL440 (2016), doi:10.1121/1.4967506.
38. E. Conan, J. Bonnel, B. Nicolas, and T. Chonavel, “Using the trapped energy ratio for source depth discrimination with a horizontal line array: Theory and experimental results,” *The Journal of the Acoustical Society of America* **142**(5), 2776–2786 (2017).
39. B. Tracey, N. Lee, and S. Turaga, “Cluster analysis and robust use of full-field models for sonar beamforming,” *The Journal of the Acoustical Society of America* **120**(5), 2635–2647 (2006).