



Published in final edited form as:

Biomed Signal Process Control. 2016 July ; 28: 19–26. doi:10.1016/j.bspc.2016.03.003.

Noise reduction in intracranial pressure signal using causal shape manifolds

Abhejit Rajagopal^{a,*}, Robert B. Hamilton^b, and Fabien Scalzo^c

^aDepartment of Electrical and Computer Engineering, University of California, Santa Barbara, USA

^bNeural Analytics Inc., Los Angeles, USA

^cDepartment of Neurology and Computer Science, University of California, Los Angeles, USA

Abstract

We present the Iterative/Causal Subspace Tracking framework (I/CST) for reducing noise in continuously monitored quasi-periodic biosignals. Signal reconstruction of the basic segments of the noisy signal (e.g. beats) is achieved by projection to a reduced space on which probabilistic tracking is performed. The attractiveness of the presented method lies in the fact that the subspace, or manifold, is learned by incorporating temporal, morphological, and signal elevation constraints, so that segment samples with similar shapes, and that are close in time and elevation, are also close in the subspace representation. Evaluation of the algorithm's effectiveness on the intracranial pressure (ICP) signal serves as a practical illustration of how it can operate in clinical conditions on routinely acquired biosignals. The reconstruction accuracy of the system is evaluated on an idealized 20-min ICP recording established from the average ICP of patients monitored for various ICP related conditions. The reconstruction accuracy of the ground truth signal is tested in presence of varying levels of additive white Gaussian noise (AWGN) and Poisson noise processes, and measures significant increases of 758% and 396% in the average signal-to-noise ratio (SNR).

Keywords

ICP; Intensive care; Traumatic brain injury; Manifold learning; Noise reduction

1. Introduction

Intracranial pressure (ICP) is defined as the pressure inside the skull and is a marker of the brain's ability to compensate for cerebral pathophysiological changes. Although intraventricular catheters and intraparenchymal sensors are widely used for ICP monitoring, their invasiveness presents considerable risk to the patient, and therefore are only used when the risks associated with a pathological increase in ICP outweighs those associated with its invasive monitoring. Clinically, ICP is a fundamental physiologic parameter that, if elevated, can lead to a pathological reduction in cerebral blood flow and possible herniation of the brain, resulting in irreversible brain damage or death if left untreated. Currently, the ICP

*Corresponding author. Tel.: +1 9492326195, abhejit@ece.ucsb.edu (A. Rajagopal).

signal is used to diagnose dangerous increases in average pressure (computed using a moving average) and helps guide treatment with recommendation based on ICP greater than 20 mmHg for more than 5 min [1]. Although the average ICP is monitored in modern clinical environments, subsequent higher-order analysis on the ICP pulse waveform often requires complex processing, expert annotations, and corrections due to egregious noise introduced during measurement from electronic equipment, electrode transients, displacement of the sensor, or even the patient shifting their posture. Such conditions make it difficult for real-time monitoring software to properly interpret ICP waveform data (Fig. 1).

The conventional approach to noise reduction is to filter the input signal. When noise is constrained to particular frequencies, such as 60 Hz tonal noise, we can retrieve most of the desired signal by applying bandstop (notch-type) filters. In the general case, when a model of the desired signal's spectral content is known, along with an estimate of the noise distribution, adaptive filtering may be used to construct a mean-square optimal filter [2]. However, the problem is more challenging when noise has a spectral density which overlaps significantly with that of the desired signal. Channel estimation is fickle in this application, because biosignals are not easily constructed from band-limited primitives. Specifically, it is increasingly difficult to identify a generic noise-floor when the relevant spectral content of the signal is not entirely known.

Another popular approach to the broadband noise problem is to estimate the original signal by assuming some stochastic mixing process. A mixture model uses knowledge of the expected degradation to estimate the most likely values for the originally transmitted signal [3]. When the number of possible source-transmitted symbols is relatively low and discretized, expectation maximization (EM) and maximum a posteriori estimation (MAP) algorithms can be used to estimate the most likely source-transmitted symbol, although an initial characterization of both the source and noise distributions is required.

In the case of ICP waveform, the number of possible pulse shapes is not finite and assigning a discrete estimate from a bank of reference signals is not an optimal solution to preserve patient-specific features. Moreover, a static characterization of the noise distribution is not always possible in clinical environments, since degradations may vary between sites and may be introduced from a combination of sources (including patient movement, sensor displacement, electronic noise, type of sensor) in varying proportion. To tackle this problem, we present a method for reducing noise in continuously monitored quasi-periodic biosignals without prior knowledge of the noise distribution. Noise is reduced by reconstructing an estimate of the original signal from a mixture of reference signals. The references are selected by searching the closest neighbors of an input sample in a reduced dimensional space. By tracking the position of consecutive samples in the subspace, causal correction transformations can be iteratively applied to the collected data stream.

There has been significant interest in modeling the morphology of the ICP waveform [4–9]. One of the main findings of those studies is to characterize the relationship between the average ICP value and the signal morphology (up to the beat level). By clustering the morphology of the waveform for various levels of ICP, as we will illustrate later in this paper, it can be observed that the ICP waveform at the beat level has shape features that vary

along with the ICP. Based on this strong correlation between the ICP value and signal morphology, we introduce a modified nonlinear subspace learning algorithm to accent this correlation by warping the learned subspace to reflect constraints about time, morphology, and average ICP. The intuition behind the algorithm is that the level of ICP implicitly constraints the possible waveform morphologies which can be used to refine the denoising process.

The structure of this paper is as follows: we first present the dataset and the ground truth used during our experimental evaluation in Sections 2.1 and 2.2. The mechanics of the Iterative/Causal Subspace Tracking (I/CST) is presented in Section 2.3, ending with the formulation of a simple I/CST implementation (Algorithm 1). The performance of this implementation is evaluated for ICP signals in Section 2.4 and the results presented in Section 3. This work results in a noise reduction framework capable of real-time performance, which enables a more reliable analysis of continuous waveform characteristics, as discussed in Section 4.

2. Methods

2.1. Data source

The dataset originated from the University of California, Los Angeles (UCLA) Medical Center, with approval from the institutional review board (IRB) for use in this study. This is a retrospective study on patients who were being treated for various intracranial pressure related conditions including idiopathic intracranial hypertension, Chiari syndrome, and slit ventricle patients with clamped shunts. A total of 60 patients were considered for this study and their ICP and electrocardiogram (ECG) signals were recorded continuously. ICP was sampled continuously at 400 Hz using an Codman intraparenchymal microsensor (Codman and Schurtleff, Raynaud, MA) placed in the right frontal lobe. An expert researcher retrospectively identified intracranial hypertension (IH) episodes and annotated the time of the elevation onset, elevation plateau, and invasive cerebrospinal fluid drainage in each patient recording. Within our cohort, 30 patients did not present any IH episodes and were excluded from the study. An additional 5 patients were excluded due to signal drop and artifacts that did not let the expert identify IH episodes with a high level of confidence. A total of 70 IH episodes were extracted from the ICP signal of the remaining 25 patients. Each segment included 20-min of data, capturing the transition from a state of normal (0–20 mmHg) to elevated ICP (>20 mmHg). The segments were time-aligned such that they contain 15 min of data before the plateau and 5 min after.

2.2. Pre-processing and ground truth

Individual ICP pulses were first extracted from each 20-min recording using a correlation of ICP with R-wave peaks in the ECG signal [10]. Because this method is dependent only locally on the R-wave peaks, the segmentation is sufficiently accurate and largely invariant to heart-rate variability [11]. The extracted pulses were distilled into 3 variables: (1) amplitude and length normalized vector containing pulsatile information, (2) mean value of the original pulse, and (3) starting time-index of the pulse relative to the elevation plateau.

The normalization is necessary because pulses extracted by the segmentation may differ in length, and require a consistent size to facilitate morphological comparisons of the form $\|x_j - x_i\|$. As such, each segmented ICP pulse s_j was normalized to a fixed-length vector $x_j \in \mathcal{R}^n$ taking values between 0 and 1, and placed into a matrix of vector samples $X \in \mathcal{R}^{m \times n}$. If the length of the pulse s_j was larger than the chosen length $n = 500$, the extra values at the end of the pulse were ignored. If the pulse was smaller than n , the last value $s_j[\text{end}]$ was repeated to fill the vector. As written formally,

$$x'_i(j=1, \dots, n) = \begin{cases} s_i[j] & \text{if } j \leq n \\ s_i[\text{end}] & \text{otherwise} \end{cases} \quad (1)$$

$$x_i(j=1, \dots, n) = \frac{x'_i - \min(x'_i)}{\max(x'_i) - \min(x'_i)} \quad (2)$$

An idealized ICP signal is then generated by accumulating similar beats and computing their average. The notion of similarity is subject to change here and is largely dependent on the intended application. In our tests, for example, beats are clustered based on their relative-time index (i.e. by binning the pulses starting within every 3 s interval). This strategy, however simple, preserves a notion of morpho-temporal locality in the ICP signal, a property that is exploited in the tracking algorithm (Section 2.3.2). Finally, to reduce errors due to misaligned segmentations, the beats in each cluster are first aligned by locating peaks in their cross-correlations.

The resulting average signal \bar{X} (Fig. 2), which still holds pulsatile information, is the “ground truth” that will further be used in our experiments. As it has been shown in previous works [9,12], the average shape of the ICP waveform is related to the ICP elevation. Pulses corresponding to normal ICP tend to exhibit three peaks (in blue), while higher ICP ones generally tend to become unimodal (in red).

For purposes of testing our algorithm, we generate two such ICP signals using similarly-sized mutually exclusive subsets of the extracted beats; that is, \bar{X}_1, \bar{X}_2 such that $|\bar{X}_1| \approx |\bar{X}_2|$, $|\bar{X}_1| \neq |\bar{X}_2|$, $\bar{X}_1 \cup \bar{X}_2 = \bar{X}$, and $\bar{X}_1 \cap \bar{X}_2 = \emptyset$. One signal (\bar{X}_1) is used for training a subspace model, and one signal (\bar{X}_2) is used for testing the tracking algorithm. We will subsequently refer to these two signals as the “reference” and “ground truth” respectively.

2.3. Iterative Causal Subspace Tracking (I/CST)

The primary objective of the proposed framework is to learn a manifold (i.e. subspace representation of the reference data) that is represented as a graph and on which consecutive noisy pulses can be projected and denoised continuously. A wide variety of manifold learning techniques exist in the literature (such as ISOMAP [13], Laplacian eigenmap [14], locally linear embedding [15]) and could be suitable for pulsatile data such as ICP. A key element of the proposed framework is to constrain the learning of the manifold to take into account the average ICP of the pulses and their relative position in time, as it can be seen in

Fig. 2 to be related to the overall shape of the pulse. By doing so, we aim that pulses with similar shapes and ICP, and close in time will also be close to each other in the manifold representation. Because the manifold will be used to track pulses over time, it is important that similar shapes of pulses remain close to each other for optimal reconstruction. The use of a graph representation allows the geodesic distance to be computed (i.e. distance in the graph, as opposed to Euclidean distance), from which variations between signals can be quantified more easily.

During training, the proposed Iterative Causal Subspace Tracking (I/CST) relies on a subspace learning algorithm (Section 2.3.1) followed by the construction of a graph defined on that space. Once learned, the graph manifold can then be used by the tracking algorithm (Section 2.3.2) to iteratively project successive noisy pulses on the graph, in order to reconstruct their mostly likely shape in the original input space.

2.3.1. Subspace learning—The subspace (Fig. 3) is obtained using a kernel discriminant analysis (KDA) of reference pulses, which is solved using a spectral regression (SR) framework [16,17]. The goal of SR-KDA in this case is to find a regression model which leads to similar subspace projections $y_j \in Y$ for input data samples (i.e. pulses) $x_j \in X$ that are morphologically similar.

SR-KDA utilizes a graph representation of the data where each vertex represents a data point. An affinity matrix $W \in \mathcal{R}^{m \times m}$ is used to represent the graph and associates a similarity weight $W_{i,j}$ to each edge $\{i, j\}$; given a set of m samples $x_{i=1, \dots, m}$. A graph embedding technique is used to represent each vertex of the graph as a vector $s_j \in \mathcal{S}$ that preserves similarities between the vertex pairs, where similarity is measured by the edge weight.

To obtain an optimal graph embedding, the objective is to ensure that samples that are close to each other in the graph are also close in the subspace representation. This can be achieved by minimizing the following measure ε :

$$\varepsilon = \sum_{i,j=1}^m (s_i - s_j)^2 \cdot W_{i,j} \quad (3)$$

$$= 2S^T L S \quad (4)$$

where $L = DW$ is the graph Laplacian and D is a diagonal matrix whose entries are column sums of W . The optimal S can be obtained by finding the largest k generalized eigenvectors λ of the eigen-problem:

$$WS = \lambda DS \quad (5)$$

Once the eigenvectors λ are computed, the embedding S of the data can be used as labels, and the regression problem solved as a standard ridge regression:

$$\arg \min_a \sum_{i=1}^m (a^T x_i - s_i)^2 + \alpha \sum_{i=1}^m a_i^2 \quad (6)$$

As mentioned earlier, the proposed framework incorporates constraints on the learning of the subspace to take into account the average ICP of the pulses and their relative-time position. This is done via the construction the W matrix so that the weights reflect a weighted distance between the respective pulse waveform p , ICP c , and time index t . That is,

$$G_1(i, j) = w_p p + w_c c + w_t t \quad (7)$$

$$G_2(i, j) = \begin{cases} 1 & \text{if } G_1(i, j) \in \text{knn}(G_1(i, :), k) \\ 0 & \text{otherwise} \end{cases} \quad (8)$$

$$W = G_1 \odot G_2 \quad (9)$$

where w_p , w_c , and w_t represent the weight associated with the different input modalities respectively, \odot represents the element-wise multiplication of matrices, and the value $k = 5$ was chosen empirically. Multiplication by the mask G_2 constrains the association of a pulse on the manifold to its k -nearest neighbors in the mixed-modality input space.

2.3.2. Tracking on the manifold—A requirement of the tracking framework is to be able to handle various levels of noise. When the noise envelope is small, the locality-preserving properties of the SR-KDA embedding ensure a projection to a reasonably small hyper-volume around the expected point. Conversely, when the noise envelope is large there is no guarantee of locality in the graph. Since our framework should be independent of the noise distribution (for applicability in clinical scenarios), a method of constraining the projection of consecutive pulses to nearby locations is necessary.

Drawing from the observation that consecutive ICP pulses are likely to exhibit similar shapes, we propose to use the trajectory of consecutive samples on the manifold as a general constraint for the denoising process. The trajectory is obtained by projecting consecutive samples into the learned subspace. A sequential tracking is then applied to estimate the most likely coordinates of the successive samples in the subspace. The coordinates are then reconstructed back to the input space using an inverse mapping to produce the denoised waveform. In essence, the procedure achieves complex non-linear predictions by employing simple prediction algorithms in the reproducing kernel Hilbert space.

Of all possible prediction algorithms, perhaps the most simple and understandable is k -nearest neighbor regression. In this algorithm, a value is constrained to the average of the values of its k -nearest neighbors. In the case of ICP signals, this translates to the average of similar waveforms, where the notion of similarity is measured using distance in the computed graph. This procedure is formalized and adapted to the I/CST framework as the RajagopalScalzo01 (RS01) denoising algorithm and is outlined below (Algorithm 1). Specifically, RS01 uses an inversely weighted reconstruction to estimate the expected signal waveform.

RS01 is designed to improve the signal quality of continuous pulsatile signals (such as ICP) existing in R^N by exploiting the locality-preserving properties of the SR-KDA embedding in the subspace Y^3 . The algorithm uses graph searching to find the k -nearest neighbors of noisy samples in a provided subspace and computes an estimate of the the expected signal using a mixture of waveforms. In order to enforce locality, the k -nearest neighbors of the previous time-sample are used rather than those corresponding to the current time-sample, thereby constraining the projection to nearby locations (Fig. 4).

RS01 also computes an error signal indicating the level of confidence in the real-time sample based on past input. The error signal is useful to identify large fluctuations in noise, as well as sudden deviations in signal properties (e.g. morphology).

The pseudocode is given below

Algorithm 1

RajagopalScalzo01 tracking algorithm

```

1: function RS01( $x_{t-1}$ ,  $x_t$ ,  $\lambda$ ,  $M$ ,  $k$ )
   let  $M$  be the manifold (training set in  $Y^m$ )
   // Project samples to subspace
2: [ $y_{t-1}$ ,  $y_t$ ]  $\leftarrow \lambda \cdot [x_{t-1}, x_t]$ 
   // Extract  $k$ -nearest neighbors in the manifold
3:  $\hat{y}_{t-1} \leftarrow knn(y_{t-1}, M, k)$ 
   // Compute distance in the graph
4: for  $i = 1$  to  $k$  do
5:    $\hat{g}_t(i) \leftarrow geodesicDist(y_t, \hat{y}_{t-1}(i), M)$ 
6: end for
   // Reconstruct the estimate
7:  $\hat{x}_t \leftarrow \sum_i \hat{y}_{t-1}(i) \hat{g}_t^{-1}(i) / \sum_i \hat{g}_t^{-1}(i)$ 
8:  $error \leftarrow \|x_t - \hat{x}_t\|$ 
9: end function

```

2.4. Experiments

To demonstrate the effectiveness of the I/CST framework, we measure the ability to effectively track input waveforms that have been degraded by various levels of noise. In

particular, we have focused on additive white-Gaussian and Poisson noise, although the method is equally valid for artifact detection via the error signal (discussed in Section 4).

Our evaluation strategy compares the signal-to-noise ratio (SNR) of input waveforms (baseline) to those produced by various denoising kernels. The goal of these kernels is to remove the noise envelope from the degraded signal and return the original pressure signal. In this context, noise is defined as any deviation from the true waveform, and is typically reported by magnitude (e.g. 2-norm). The SNR is calculated on a beat-by-beat basis as follows.

$$SNR = \frac{\text{SignalPower}}{\text{NoisePower}} \propto \frac{\text{SignalMagnitude}}{\text{NoiseMagnitude}} \quad (10)$$

$$SNR_{\text{baseline}} = \frac{\|x(t)_{\text{true}}\|_2}{\|x(t)_{\text{true}} - x(t)_{\text{noisy}}\|_2}$$

$$SNR_{\text{RS01}} = \frac{\|x(t)_{\text{true}}\|_2}{\|x(t)_{\text{true}} - x(t)_{\text{recovered}}\|_2} \quad (11)$$

In our evaluation we used the ground truth data generated from the 70 IH episode dataset (Sections 2.1 and 2.2) to synthesize noisy signal streams representing various levels of degradation. AWGN is typically parametrized by the parameters $\mu = 0$ and σ , which represent the mean and standard deviation of the distribution. As such, a zero-mean Gaussian distribution was sampled and its variance scaled relative to the normalized ICP pulse amplitude, $\max x_{\text{ICP}}$. This noise profile was applied additively to the full length of the ICP signal to generate noisy testing data. Similarly, Poisson noise is typically parametrized by the parameter λ which represents both the mean and variance. To generate different levels of Poisson-noise, the original ICP stream was used as the mean, while the variance was similarly scaled relative to the normalized ICP pulse amplitude.

In our tests, the RS01 algorithm is compared against various Gaussian low-pass filters (LPF). These generic convolutional filters (kernels) were constructed by generating several ideal zero-mean Gaussian distributions with variances corresponding to different levels of noise. The value of variance was scaled relative to the normalized ICP amplitude, and the length of the kernel was chosen to be sufficiently large to avoid edge effects during convolution. In our tests, we compare against three such generic kernels, termed LPF1, LPF2, LPF3, corresponding to variances of 5%, 10%, and 20% of the normalized ICP amplitude.

In short, our testing procedure is comprised of four steps: (1) degrade the original ground truth signal with the selected noise profile, (2) apply the selected denoising kernel, (3)

measure the SNR of each beat, (4) average the SNR over the entire ICP signal. Our results are presented in the next section.

3. Results

Although the generic filters do provide smoothing, we find that even the simplest I/CST implementation provides a significantly more desirable result. Over all our experiments, RS01 performed consistently well with an average SNR improvement of 758% and 396% for AWGN and Poisson noise respectively. The average SNR was computed from experiments with noise variances uniformly distributed on 5% to 35% of the normalized beat amplitude, and is showcased in Table 1.

One attractive feature of the RS01 algorithm is that it does not require knowledge of the noise profile to be effective. Although such information can be useful to effectively clean up the signal, the typical convolutional approach is limited for two reasons: (1) the size of the averaging kernel is not easily determined directly from the input data, making it impractical without proper calibration or channel-estimation protocols, and (2) convolution in the time domain corresponds to a multiplication in the frequency domain, so typical Gaussian filters will effectively mask potentially-useful high-frequency information. Our results demonstrate this via a dramatic increase in the SNR, particularly at higher noise powers (Fig. 5).

Although a given kernel may outperform RS01 at a particular noise level, RS01 significantly outperforms typical low-pass filters when evaluated across various noise magnitudes. This result is elucidated by the variable SNR performance of any-single LP filter. From these tests, it is evident that no single generic filter can be unanimously selected for denoising since its performance depends on the level of degradation, which is essentially unpredictable in clinical scenarios. In contrary, RS01 pays a small penalty in initially but performs consistently well across several noise levels, making it a practical tool for clinical data collection and analysis.

It should be noted, however, that RS01 underperforms at very low noise powers (less than 3% of the signal magnitude). This is because signals are shifted slightly from their appropriate location by the local averaging effect. Even so, even at 5% relative noise amplitude, RS01 provides significant improvement to the signal quality.

The power of RS01 is evident both quantitatively from the SNR computation, and qualitatively by visualizing the denoised waveform (Fig. 6). In our tests, RS01 consistently provides a smooth signal, while various time-domain filters do not guarantee such a criteria particularly at the boundaries of each pulse-beat. Furthermore, generic filters seem to expose several local minima and maxima, which we anticipate will reduce the accuracy of subsequent higher-order processing algorithms, such as MOCAIP, while RS01 seems to match both the number, location, and amplitude of salient waveform peaks. In this respect, the RS01 kernel provides a result even better than that suggested by the 2-norm errors.

4. Discussion

Generally, analyzing transformations of signals casts the problem of recovering source-transmitted information as a class identification problem, rather than a time-domain denoising problem. A range of classification methods exist, typically in the form of vector quantization (MAP, ML, SVM) or tree-based (Bayesian neural network) decision rules. These schemes, when applied to waveforms which map to a discrete symbol distribution, have demonstrated incredible improvements in both the symbol-detection accuracy and effective signal-to-noise ratio in the presence of degradations [18–20]. Unfortunately, biosignals generally do not map to a discrete distribution and such estimates are not easily extended to the high-dimensional large multi-class problem. Organic signals, such as ICP, tend to exhibit a range of features that are locally temporally correlated, but vary continuously with the waveform [21]. In this light, our goal was to develop a robust information-processing system for analog biosignals.

In contrast to previous works [22,9], our method captures the continuously varying characteristics of ICP waveforms. This was accomplished by characterizing the morphology of ICP waveforms via clustering, warping the subspace using continuous valued statistics (DC value and expert annotations), and tracking the progression of waveforms within the proposed I/CST framework. In developing this methodology, we discovered a way to reconstruct a likely estimate of the original signal in noisy scenarios using a mixture of waveforms (Algorithm 1). By searching the k -nearest neighbors of a sample in the learned subspace, we effectively constrain and denoise the analog estimate. The benefit of such a scheme was demonstrated by significant increases in the average and peak SNR of a 20-min ICP recording.

The resilience of RS01 to high-levels of noise (particularly AWGN) can perhaps be attributed to the averaging characteristic of the chosen mapping metric (Euclidean distance). Because deviations at any position in the signal contribute to the error during morphological comparisons, the error can be minimized by choosing waveforms that match the natural oscillations present in the noisy signal. However, the averaging characteristic can also be detrimental when the original waveform is largely smooth with few local oscillations. In this respect, the Euclidean metric is somewhat naive. In particular, although the peak and mean SNR are higher, our implementation suffers from errors primarily in regions where the source transmitted signal is unimodal (after the IH elevation plateau). This is evident by tracking the SNR of each time interval. This issue could be remediated by defining an alternative morphological comparison metric, which can be easily supplied to both RS01 and the I/CST framework.

More broadly, I/CST provides a convenient framework to improve the quality and analysis of measured biosignals by providing a mechanism for triggered analysis and qualification of noisy samples. The benefit of this solution is that new samples can be projected in real-time, and subsequent analysis can be performed in parallel to qualify the trajectory and properties of projected nodes. One of the main insights in this design is that the trajectory (i.e. a directed path on a graph) of consecutive samples can be used to augment tracking and analysis algorithms. Furthermore, I/CST's ability to operate hierarchically on N -dimensional

observations in real-time, as opposed to retrospectively, provides a robust platform for developing automated adaptive software systems, including control and learning systems which rely on programmed routines but require live observation-based triggers.

One pitfall of our presented implementation is that it is heavily reliant on the training period. A careful characterization of the ICP signal, typically done via clustering, is required to accurately reconstruct signals using their projections. In particular, we find that the frame-width is considerably important; coarse divisions may detract from tracking ability, while fine divisions may detract from accurate quantization solutions. In our work, we chose to cluster ICP signals based on their distance from an expert annotation (IH plateau) since we hypothesized this to be a relevant marker; so, in this work the window size corresponds to time intervals, empirically chosen to be 3.75 s and producing a constellation of 320 symbols. Such a decision must be made a posteriori, and thus requires domain-specific knowledge.

One way to generalize the I/CST approach to generic signals (i.e. processing without domain-specific knowledge) is to first employ morphological clustering and subsequently learn the subspace with continuous annotations. The morphological clustering can be domain-specific, but it is not necessary. A method of N -dimensional k -means clustering and regression could be applied to identify several possible distinct symbols which vary proportionally with some statistic (mean, variance, kurtosis, etc), although an appropriate clustering procedure is highly non-trivial. Obviously, the statistics used should be relevant to the waveforms being compared for optimal performance, but this criteria is not necessary for the basic operation of the algorithm.

In this sense, the I/CST framework is sufficiently general and by itself does not require any domain-specific knowledge. As such, I/CST can be easily extended to other quasi-periodic biosignals given an appropriate clustering and comparison strategy. Examples include electrocardiogram (ECG), transcranial Doppler (TCD), and generic pulsatile signals such hormone or protein sensors.

5. Conclusion

The elementary operations employed by I/CST and RS01 are not new. They are, however, employed in a novel way. In a sense, this work demonstrates that linear (or other) prediction algorithms need not be constrained to a particular domain, such as time or frequency. These domains are typically used because they make apparent certain features that aid in the discrimination of signals (such as amplitude and frequency peaks, or phase dispersion). We generalize this concept to show that signals, such as ICP, can be cast into an arbitrary discriminative domain where previously-developed elementary algorithms are still effective. In particular, by reducing the dimensionality of the subspace where signals are projected, the operation of such algorithms is considerably improved (both in complexity and accuracy) due to the low dimensionality of the search hyper-volume. This result can be exploited by medical and biological analysis software for the purpose of state-prediction and active denoising.

References

1. Faul M, Wald MM, Rutland-Brown W, Sullivent EE, Sattin RW. Using a cost-benefit analysis to estimate outcomes of a clinical treatment guideline: testing the brain trauma foundation guidelines for the treatment of severe traumatic brain injury. *J. Trauma Acute Care Surg.* 2007; 63(6):1271–1278.
2. Shynk JJ, et al. Frequency-domain and multirate adaptive filtering. *IEEE Signal Process. Mag.* 1992; 9(1):14–37.
3. Lu L, Chin K, Ghoshal A, Renals S. Noise compensation for subspace Gaussian mixture models. *INTERSPEECH.* 2012
4. Cardoso ER, Rowan JO, Galbraith S. Analysis of the cerebrospinal fluid pulse wave in intracranial pressure. *J. Neurosurg.* 1983 Nov.59:817–821. [PubMed: 6619934]
5. Dubin MJ, Magram G, Prasad AK. Intracranial pressure waveform analysis: computation of pressure transmission and waveform shape indicators. *Neurol. Res.* 1998 Sep.20:533–541. [PubMed: 9713845]
6. Czosnyka M, Smielewski P, Timofeev I, Lavinio A, Guazzo E, Hutchinson P, Pickard JD. Intracranial pressure: more than a number. *Neurosurg. Focus.* 2007; 22(5):E10.
7. Heldt, T., Verghese, G. Modeling the morphology of the intracranial pressure waveform; *IEEE Engineering in Medicine and Biology Conference (EMBC);* 2013. p. 1
8. Scalzo F, Xu P, Asgari S, Bergsneider M, Hu X. Regression analysis for peak designation in pulsatile pressure signals. *Med. Biol. Eng. Comput.* 2009 Sep.47:967–977. [PubMed: 19578916]
9. Scalzo F, Asgari S, Kim S, Bergsneider M, Hu X. Bayesian tracking of intracranial pressure signal morphology. *Artif. Intell. Med.* 2012; 54(2):115–123. [PubMed: 21968205]
10. Choi S, Adnane M, Lee G-J, Jang H, Jiang Z, Park H-K. Development of ECG beat segmentation method by combining lowpass filter and irregular R-R interval checkup strategy. *Expert Syst. Appl.* 2010; 37(7):5208–5218.
11. Scalzo F, Asgari S, Kim S, Bergsneider M, Hu X. Robust peak recognition in intracranial pressure signals. *Biomed. Eng. Online.* 2010; 9(1):61. [PubMed: 20959014]
12. Scalzo F, Liebeskind D, Hu X. Reducing false intracranial pressure alarms using morphological waveform features. *IEEE Trans. Biomed. Eng.* 2013; 60(1):235–239. [PubMed: 22851230]
13. Tenenbaum JB, De Silva V, Langford JC. A global geometric framework for nonlinear dimensionality reduction. *Science.* 2000; 290(5500):2319–2323. [PubMed: 11125149]
14. Belkin M, Niyogi P. Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput.* 2003; 15(6):1373–1396.
15. Roweis ST, Saul LK. Nonlinear dimensionality reduction by locally linear embedding. *Science.* 2000; 290(5500):2323–2326. [PubMed: 11125150]
16. Cai, D., He, X., Han, J. 2007 Seventh IEEE International Conference on Data Mining, ICDM 2007. IEEE; 2007. Efficient kernel discriminant analysis via spectral regression; p. 427-432.
17. Cai D, He X, Han J. Speed up kernel discriminant analysis. *VLDB J.* 2011; 20(1):21–33.
18. Proakis JG, Salehi M, Zhou N, Li X. *Communication Systems Engineering, 1*, Prentice Hall. Englewood Cliffs. 1994
19. Prasad, GK., Sahambi, J. Conference on Convergent Technologies for the Asia-Pacific Region, TENCON 2003. Vol. 1. IEEE; 2003. Classification of ECG arrhythmias using multi-resolution analysis and neural networks; p. 227-231.
20. Allaire, S., Kim, JJ., Breen, SL., Jaffray, DA., Pekar, V. 2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition Workshops, CVPRW'08. IEEE; 2008. Full orientation invariance and improved feature selectivity of 3D SIFT with application to medical image analysis; p. 1-8.
21. Neagoe, V-E., Iatan, I-F., Grunwald, S. AMIA Annual Symposium Proceedings. Vol. 2003. American Medical Informatics Association; 2003. A neuro-fuzzy approach to classification of ECG signals for ischemic heart disease diagnosis; p. 494

22. Hu X, Xu P, Scalzo F, Vespa P, Bergsneider M. Morphological clustering and analysis of continuous intracranial pressure. *IEEE Trans. Biomed. Eng.* 2009; 56(3):696–705. [PubMed: 19272879]

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

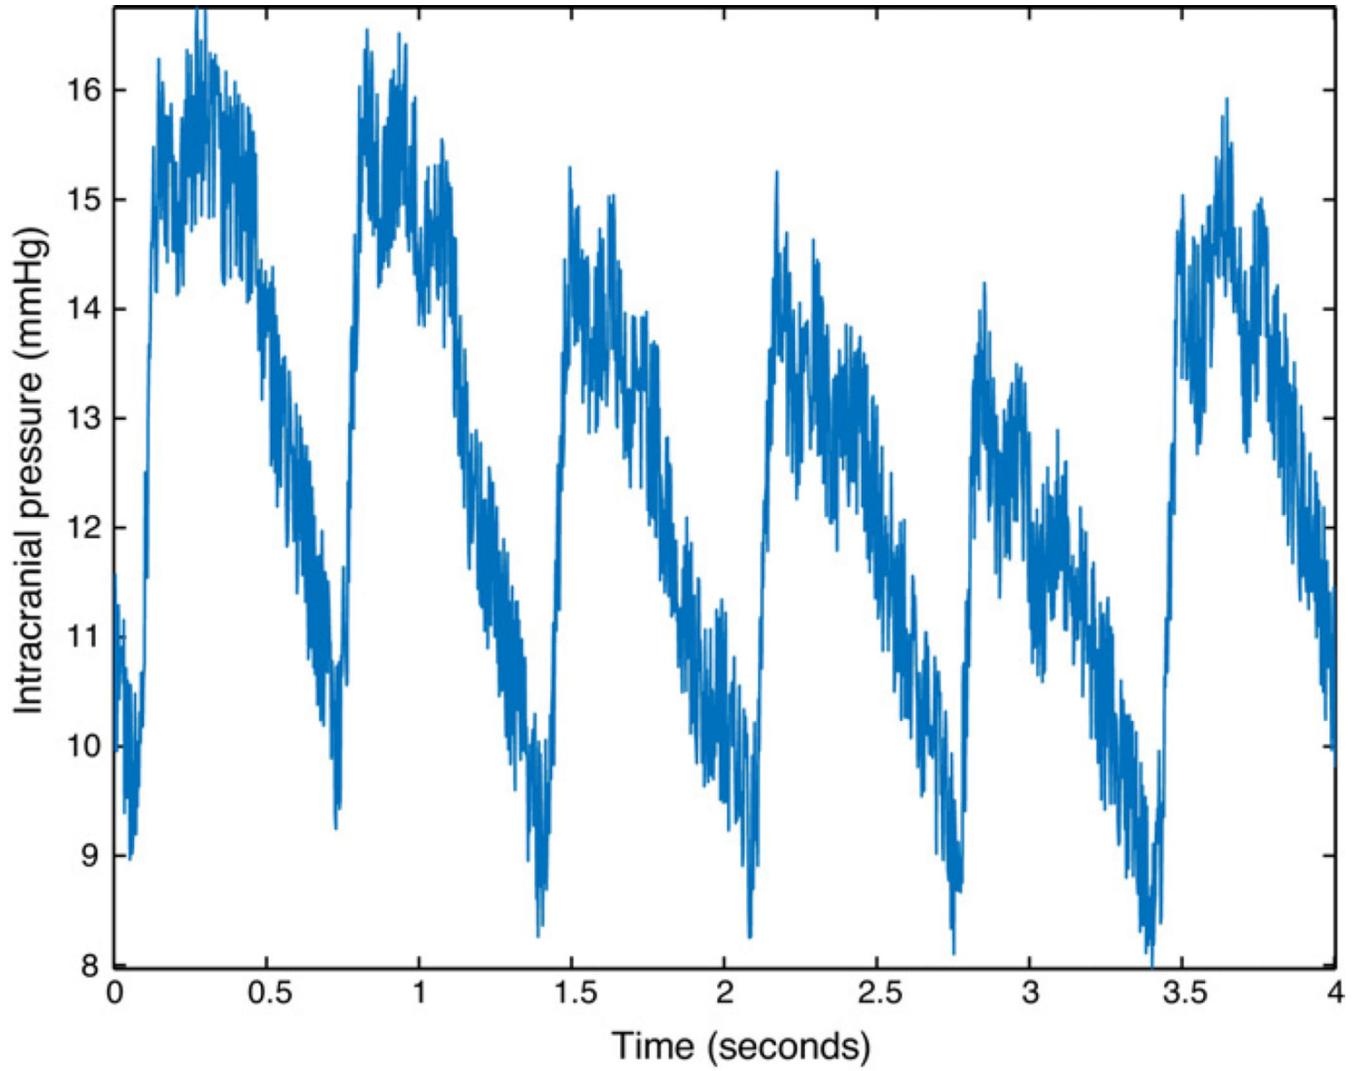


Fig. 1. Illustration of a typical intracranial pressure signal (ICP) recorded in clinical conditions. ICP typically exhibits significant noise on its envelope that challenges its morphological analysis.

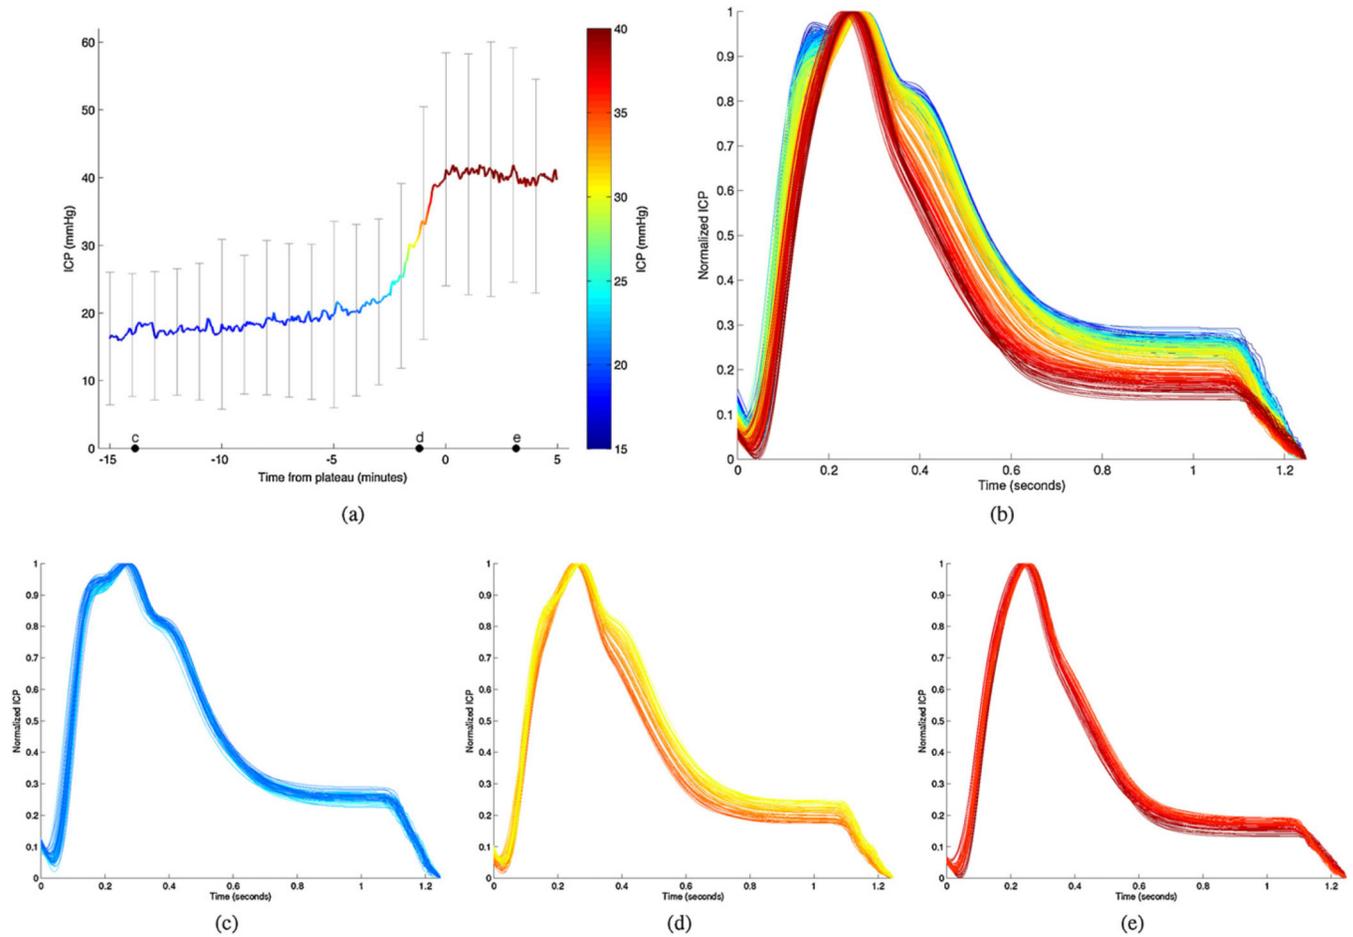


Fig. 2.

(a) The average and standard deviation of pulsatile ICP clusters at each time point relative to hypertension onset. Panel (b) depicts the morphology of time-localized ICP clusters. Three groups of such clusters are illustrated for the baseline (c), transition (d), and hypertensive regions (e). (For interpretation of the references to color in the text, the reader is referred to the web version of this article.)

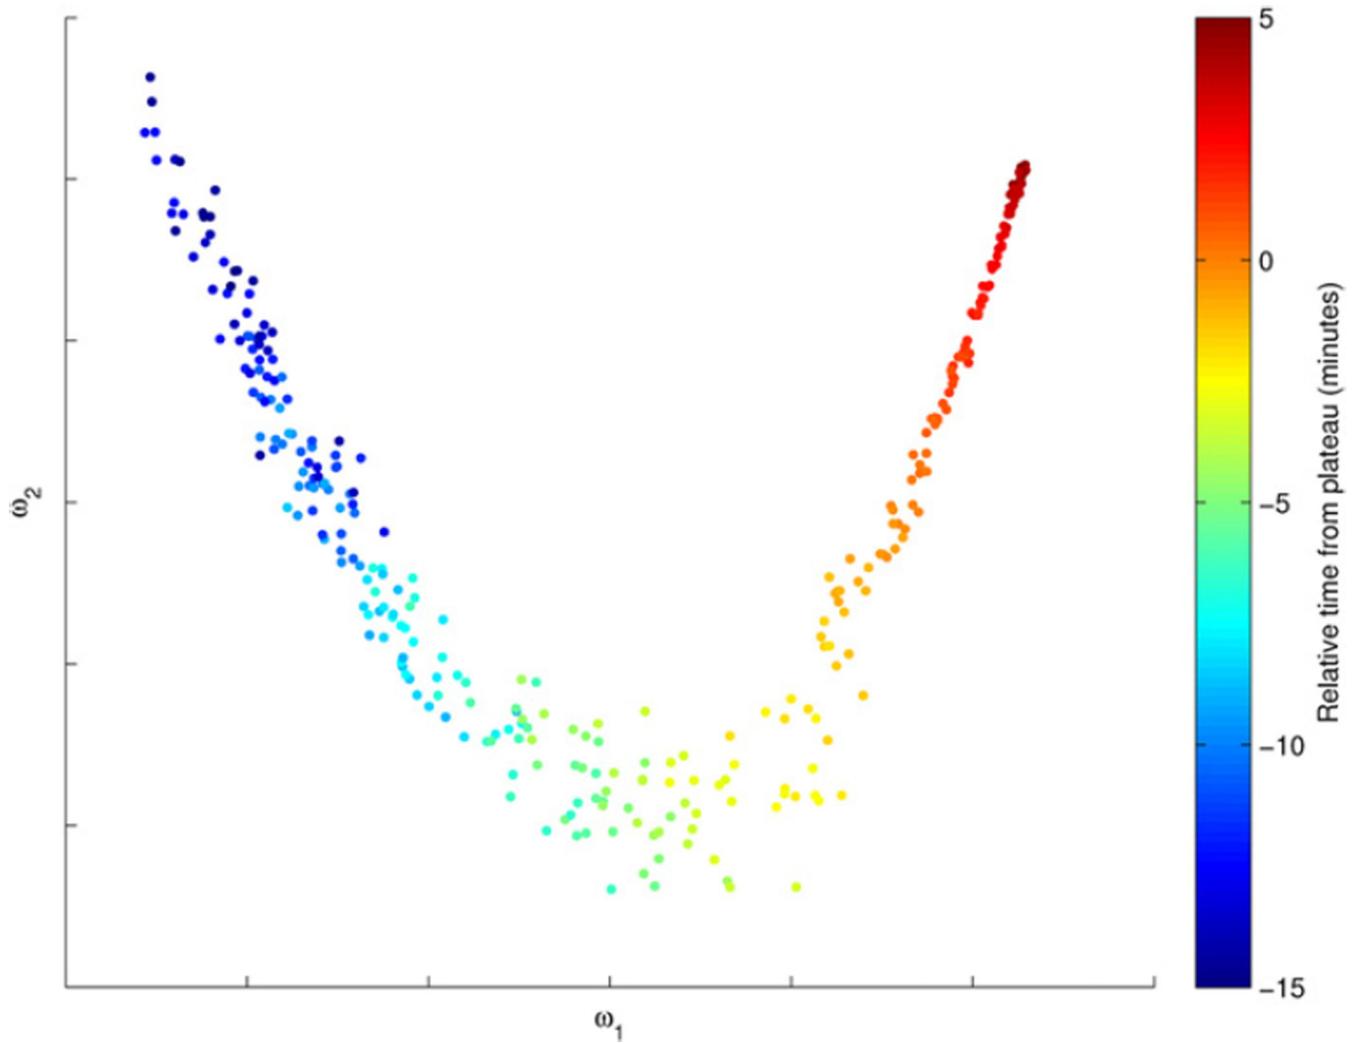


Fig. 3. Projection of a set of ICP beats onto the first two dimensions (ω_1 and ω_2) of the manifold. Each dot represents a 3-s cluster extracted from the training set. The transition from normal (blue) to elevated (red) ICP is depicted by a colormap that represents the temporal-index relative to the elevation plateau. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

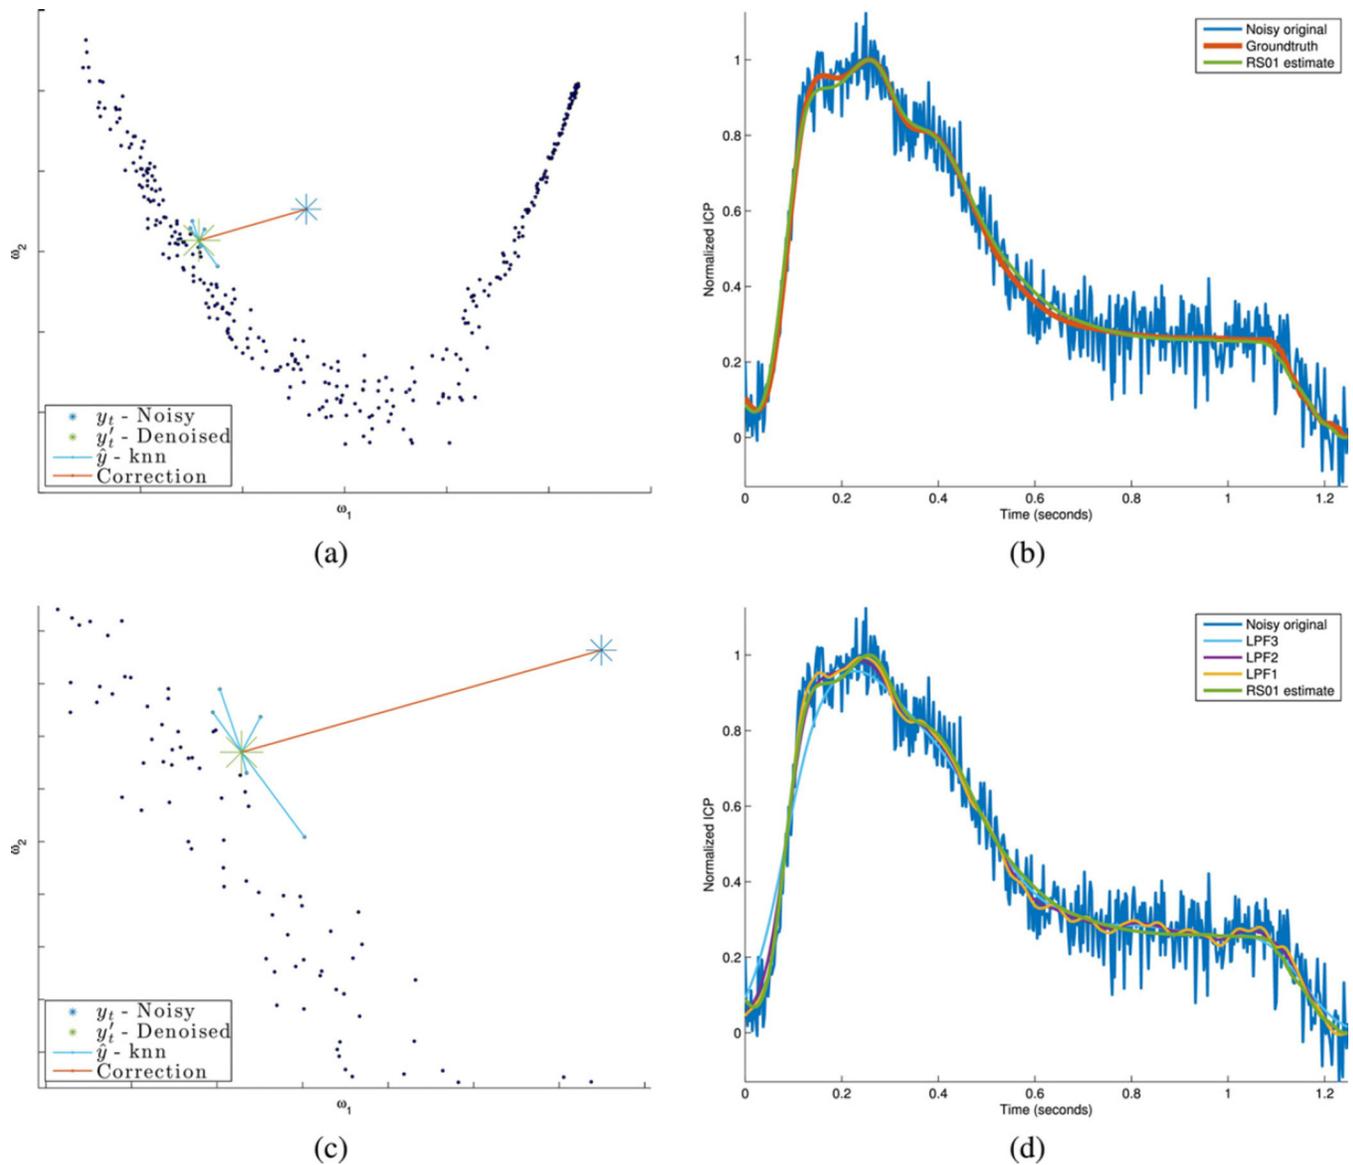


Fig. 4. RS01 operates by projecting the current sample onto a subspace defined by a training constellation (a), and constraining this projection to a mixture of the k -nearest neighbors of the previous sample (b). The final estimate is a smooth signal (c) which demonstrates significant improvements over generic filters (d) in terms of smoothness, mean-square error, and distinguishable morphology such as the number and location of local maxima and minima.

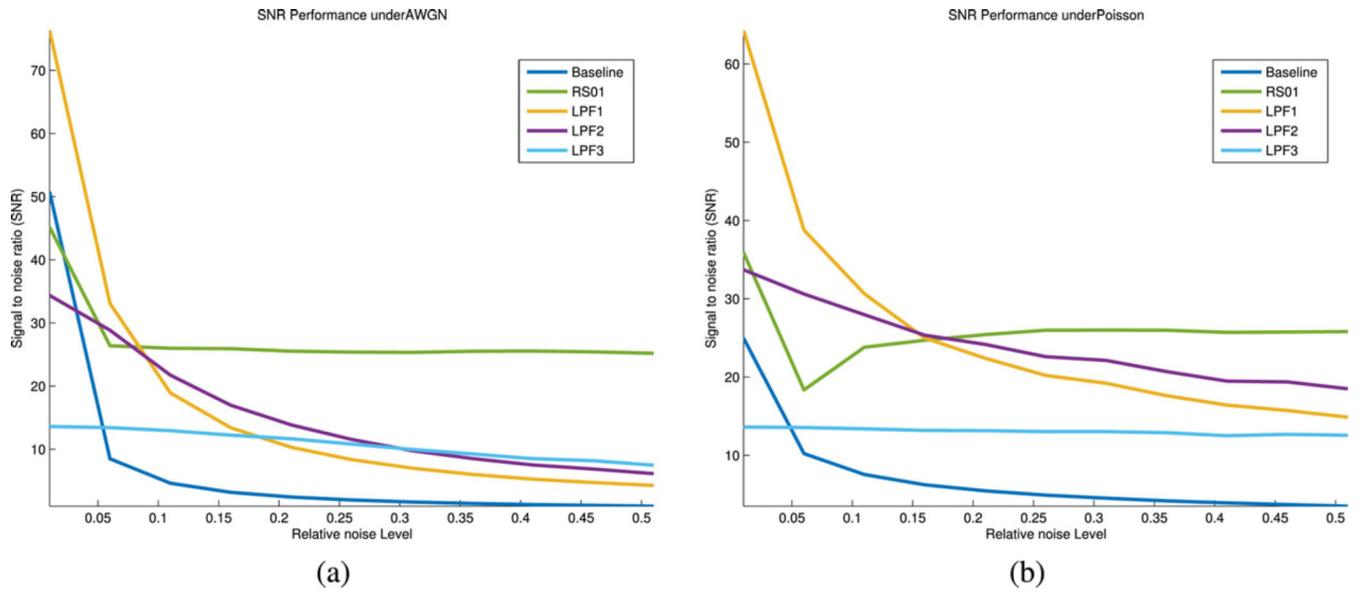


Fig. 5. SNR at various AWGN and Poisson-noise magnitudes.

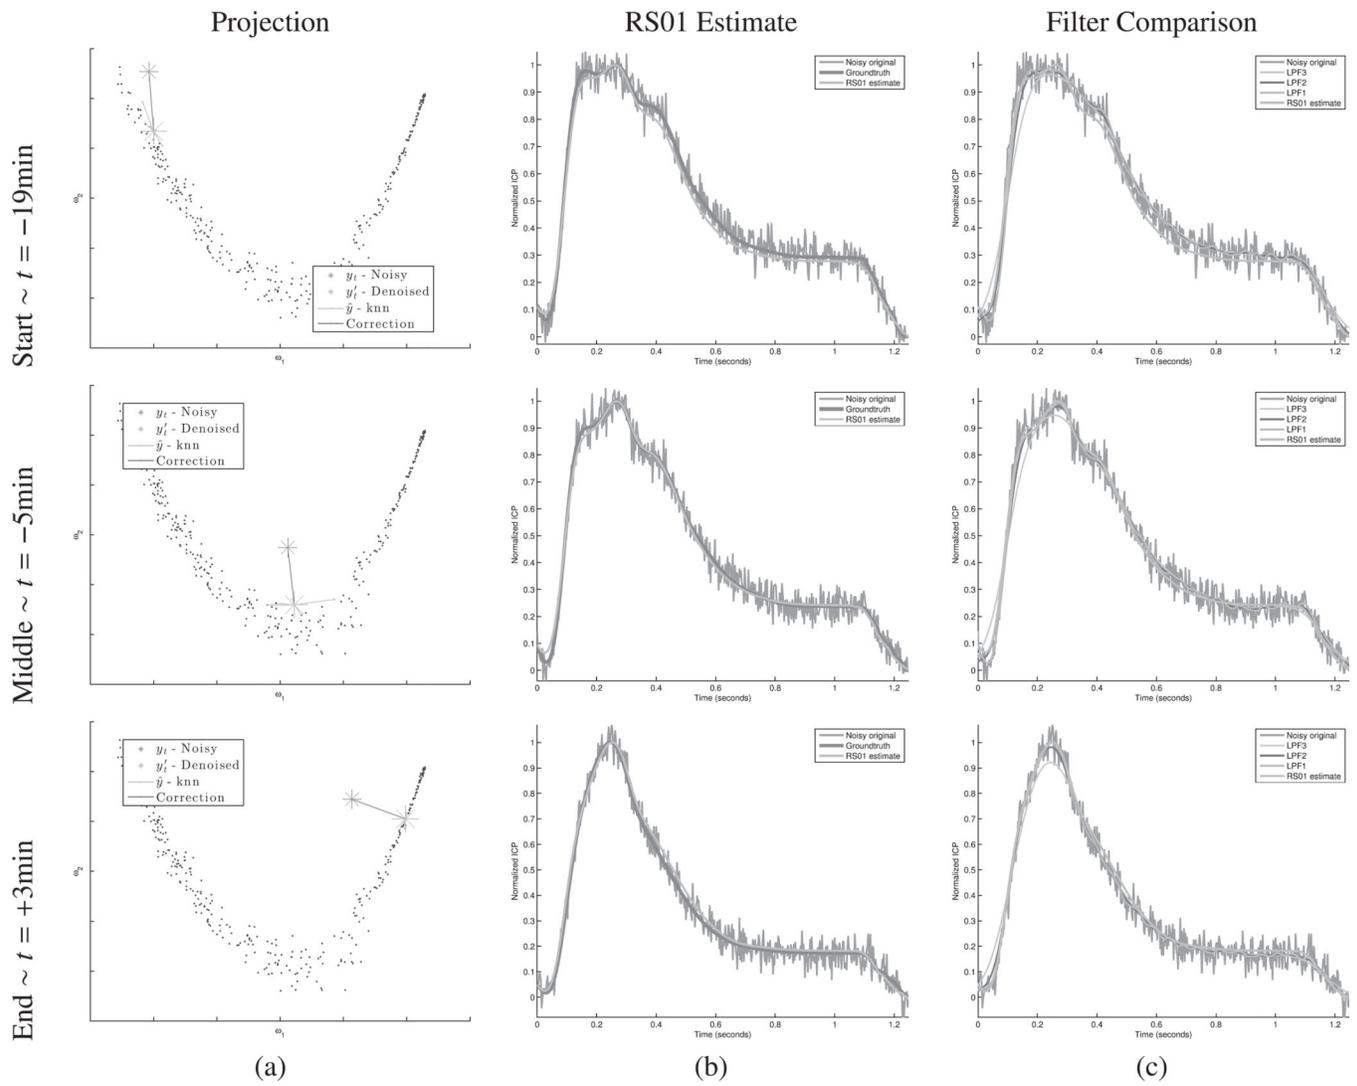


Fig. 6. Denoising of an idealized ICP stream corrupted by AWGN noise of variance $\sigma = 4\%$. (a) Signals are projected to the subspace and (b) denoised using the RS01 algorithm using geodesic constraints. (c) The estimate is compared to various generic filters.

Table 1

The average signal-to-noise ratio (SNR) across tests, with variance ranging from 5% to 35% of the normalized beat amplitude. RS01 outperforms generic filters as the noise magnitude increases.

Kernel	AWGN				Poisson noise			
	5%	15%	25%	35%	5%	15%	25%	35%
Baseline	10.1539	3.3948	2.0354	2.0337	11.1322	6.4324	4.9938	4.9791
RS01	28.3347	25.8013	25.5826	25.8380	19.0380	24.2369	25.8184	25.9031
LPF1	38.3530	14.2530	8.5328	8.6612	41.0770	26.2788	20.8863	20.7245
LPF2	30.1515	17.9025	11.7249	11.8329	30.7207	26.0763	23.2576	22.9546
LPF3	13.3555	12.4608	10.9168	10.9168	13.4938	13.2455	13.0712	13.0037

Bold indicates the highest SNR for each given experiment.