Published in final edited form as:

Proc IEEE Int Conf Acoust Speech Signal Process. 2016 March ; 2016: 6290–6294. doi:10.1109/ICASSP.2016.7472887.

THE INTRINSIC VALUE OF HFO FEATURES AS A BIOMARKER OF EPILEPTIC ACTIVITY

Stephen V. Gliske*,

University of Michigan, Department of Neurology, Ann Arbor, Michigan, U.S.A

William C. Stacey*,

University of Michigan, Department of Neurology, Ann Arbor, Michigan, U.S.A

Kevin R. Moon†, and

University of Michigan, Electrical Engineering and Computer Science, Ann Arbor, Michigan, U.S.A

Alfred O. Hero III†

University of Michigan, Electrical Engineering and Computer Science, Ann Arbor, Michigan, U.S.A

Abstract

High frequency oscillations (HFOs) are a promising biomarker of epileptic brain tissue and activity. HFOs additionally serve as a prototypical example of challenges in the analysis of discrete events in high-temporal resolution, intracranial EEG data. Two primary challenges are 1) dimensionality reduction, and 2) assessing feasibility of classification. Dimensionality reduction assumes that the data lie on a manifold with dimension less than that of the features space. However, previous HFO analysis have assumed a linear manifold, global across time, space (i.e. recording electrode/channel), and individual patients. Instead, we assess both a) whether linear methods are appropriate and b) the consistency of the manifold across time, space, and patients. We also estimate bounds on the Bayes classification error to quantify the distinction between two classes of HFOs (those occurring during seizures and those occurring due to other processes). This analysis provides the foundation for future clinical use of HFO features and guides the analysis for other discrete events, such as individual action potentials or multi-unit activity.

Index Terms

high-frequency oscillation; intrinsic dimension; dimensionality reduction; classification error; divergence

1. INTRODUCTION

About one third of epilepsy patients fail to obtain seizure control with available pharmaceuticals. One of the few options for these refractory patients is resective surgery—removing the portion of the brain thought to be causing the seizures. This region is denoted

*These authors were partially supported by the Doris Duke Foundation and the US National Institute of Health (NIH) grant 5-K08-NS069783-05. SG was additionally supported by an NIH Big Data to Knowledge Mentored Training Grant, 1-K01-ES026839-01A1.

†These authors were partially supported by the US National Science Foundation (NSF) under grant CCF-1217880 and a NSF Graduate Research Fellowship to KM under Grant No. F031543.

the seizure onset zone (SOZ). In some cases, determining the SOZ involves a highly invasive surgery to place electrodes on the brain's surface, followed by one to two weeks of recording and monitoring. A second invasive surgery is performed if the SOZ can be identified and safely resected. A schematic relating the implanted electrodes with the recorded data is shown in Fig 1.

A proposed biomarker to improve the localization of the SOZ are high frequency oscillations (HFOs) [1, 2]. HFOs are high frequency (about 80–300 Hz), short (< 50 ms), rare events occurring in intracranial EEG recorded at sampling rates of several kHz. Example HFO detections and a recorded seizure are shown in Fig 2.

Much of the published research on HFOs uses human identified HFOs in short (10 to 20 minute) recordings [3, 4]. These results have shown that a high HFO occurrence rate is correlated with the SOZ. However, recent work is moving towards automated identification and analysis of HFOs in long term, high resolution data, which requires advanced computational and statistical techniques [5]. Quality recordings may span 7–14 days, with over 100,000 HFO detections in several terabytes of data. Thus, the next advances are expected to come through big data analysis of HFO features, utilizing millions of recorded HFOs across as many patients as possible [6].

Relatively few research groups have analyzed HFO features in detail. The most advanced analysis computed six features of about 300,000 HFOs in nine patients and two controls, and utilized a global PCA across all channels followed by k -means clustering [7, 8, 9]. The authors implicitly assumed that the distribution of these HFO features lies on a linear manifold in feature space, and that the manifold is consistent across time and space (i.e., recording electrode).

The other most advanced analysis compared HFOs produced in the motor cortex via movement versus HFOs occurring in the SOZ, utilizing three features and a support vector machine (SVM) classifier [10]. Some differences were noted, but a more general analysis that addresses the degree to which HFOs produced by pathological activity or networks (denoted pathological HFOs or pHFOs) and HFOs produced by normal, physiological activity (denoted normal HFOs or nHFOs) are observably different has not been performed.

The goals of this work are to test the implicit assumptions previously used in HFO feature analysis. Specifically, the goals are to 1) assess the type of manifold on which HFO features lie, and 2) assess how discernible pHFOs are from nHFOs, based on their feature-space distributions.

The general outline is as follows. To assess the linearity of the HFO-feature manifold, a non-linear, local estimate of intrinsic dimension [11, 12] is compared with a linear global estimate (PCA) applied to local subsets. This approach is similar to that in [13]. The corresponding reduced dimension subspaces are individually compared across time and space using a modified Grassmann distance, building off the work in [14]. We additionally use a greedy Fisher LDA algorithm (similar to the greedy LDA in [15]) to identify a basis that maximally separates pHFOs from nHFOs. Unsupervised clustering of the subspaces is then compared with channel groups based on clinical markings of the SOZ and physical

groupings of the electrodes. Lastly, we assess the discernibility of nHFOs and pHFOs by estimating bounds on the Bayes Error using the Henze-Penrose divergence (HPD) [16, 17].

2. PATIENT POPULATION AND DATA

EEG data from adult patients with refractory epilepsy who underwent intracranial EEG monitoring were selected from the IIEG Portal [18] and from the University of Michigan. All patient data was included which met the following criteria: sampling rate of at least 5 kHz, recording time greater than one hour, data recorded with traditional intracranial electrodes, and available meta-data regarding seizure times and the resected volume or SOZ. This yielded 17 patients, (nine IIEG portal, eight U. of M.). All data were acquired with approval of local institutional review boards (IRB), and all patients consented to share their de-identified data. Of these 17 patients, 13 had recorded seizures, nine were known to have resection and obtained seizure-freedom (ILAE class I), with eight patients in common between these two categories. Because these surgeries are relatively rare, this patient population size is moderately large for analysis of intracranial EEG data.

HFOs were detected using the qHFO algorithm [5], resulting in over 1.6 million HFOs in nearly 100,000 channel-hours of 5 kHz data. Each HFO was band pass filtered between 80 to 500 Hz using an elliptical filter and 33 features were computed, including duration, peak power, mean of the Teager-Kaiser energy [19], and various spectral properties. Ictal is defined as during seizures, with seizures assumed to be five minutes long if the length was not specified in available meta-data. Interictal is defined as at least 30 minutes from the start or end of a seizure, based on [9].

3. DIMENSIONALITY REDUCTION

3.1. Consistency of local, non-linear intrinsic dimension

To assess both the linearity and local versus global nature of the HFO feature manifold, the non-linear intrinsic dimension was computed via the k -nearest neighbor (NN) based estimator in [11]. This estimator estimates the intrinsic dimension by exploiting its relationship to the total edge length of the k -NN graph. This nonlinear estimator provides a *local* estimate of intrinsic dimension which enables us to identify local variations in data manifolds. The result is an estimate of the intrinsic dimension for each given HFO, which are then averaged to obtain the mean intrinsic dimension for a given partition of the data. The consistency of the manifold is measured by comparing the distribution of intrinsic dimension across time, space and patients. The specific comparisons are: a) interictal versus ictal times, per channel, b) time variation within interictal periods, per channel, and c) comparisons between channels, integrated over time.

A variety of methods could be employed to compare the intrinsic dimension for two disjoint sets of HFOs. However, the final dimension selected will be an integer. Thus, small differences in the intrinsic dimension between two sets, no matter how statistically significant, are not meaningfully different if they mean value for each set round to the same integer.

To compare two sets of intrinsic dimension, we define a distance measure, θ_I , between collections of integers. Let the two sets of integers be A and B , and let n_i be the fraction of elements in A equal to i , and m_i be the fraction of elements in B equal to i . The probability that two elements in set A are equal is $\mathbf{n}^T \mathbf{n}$, and the probability that an element of A is equal to an element of B is $\mathbf{n}^T \mathbf{m}$. A measure of distance between A and B is how likely an element of A is equal to an element of B , normalized by the likelihoods of elements being equal another element in the same group:

$$\theta_I = \arccos \left(\frac{\mathbf{n}^T \mathbf{m}}{\sqrt{\mathbf{n}^T \mathbf{n}} \sqrt{\mathbf{m}^T \mathbf{m}}} \right), \quad (1)$$

The value θ_I is the angle between $\mathbf{n}$ and $\mathbf{m}$ and provides an easy interpretation as to the consistency of two different collections of local intrinsic dimensional. This quantity is also known as the angular distance or angular dissimilarity, and is the inverse cosine of the Ochiai-Barkman coefficient [20, 21].

We compute the θ_I -distance for three different comparisons of HFOs: 1) a comparison of each pair of 30 minute time windows on a given channel for a given patient, 2) a comparison of ictal versus interictal periods, again on a given channel for a given patient, and 3) a comparison of different channels in a given patient during interictal periods. Interictal-ictal comparisons where one set of events is less than 50 HFOs are ignored. Histograms of the distribution of θ_I for each type of comparison are shown in Fig. 3. This figure involves over 5 million interictal time bin comparisons, 163 ictal versus interictal comparisons, over 36 thousand interictal channel-channel comparisons and almost 10 thousand ictal channel-channel comparisons.

The intrinsic dimension is quite consistent across different time segments during interictal times (strong peak near 0°), but has some variance between ictal and interictal times (still peaked near 0° but the peak is wider). The dimension is less consistent across channels, with comparisons during ictal times showing small peaks at both 0° and 90° , and interictal having the largest peak at 90° , showing maximal difference. Thus we see that the intrinsic dimension varies significantly across channels, especially for interictal HFOs.

3.2. Comparison with Global Linear Intrinsic Dimension

Next we compare the k -NN intrinsic dimension estimate of Section 3.1 with a global, linear method to assess how linear and/or local the feature manifold is. The most common global, linear method of intrinsic dimension estimation and dimensionality reduction is principle component analysis (PCA), which we perform by first centering the data and then using singular value decomposition (SVD). PCA is performed multiple times for the same divisions of the HFO data as done for Fig. 3. Subsets of HFOs with less than 50 events are ignored, as these are deemed insufficient to estimate the PCA vectors in 33D. This results in PCA vectors being computed for 606 of the 1318 channels (12 patients) for interictal HFOs, 171 channels for ictal HFOs (8 patients), and 163 channels comparing ictal versus interictal

(8 patients). In patients that had multiple recording sessions, channels are counted once per each session.

To compare the k -NN (non-linear) and PCA intrinsic dimension per channel, we 1) select the number of principle components equal to the non-linear intrinsic dimension estimate, and then 2) report the fraction of the variance accounted for by that number of principle components. This is repeated for all 606 channels (interictal) and 171 channels (ictal). When using either ictal or interictal HFOs, the median fraction of variance was 99.8%, with 99%-tile of the channels being above 89.5% (interictal) and 97.1% (ictal). It is likely that that noise in the data could account for up to 10% of the variance, and thus we conclude that the feature manifolds are approximately linear over the locality of a given channel and ictal state.

3.3. Comparison between Local Manifolds

In addition to the earlier comparison of the subspace dimensionality across channels, the next step is to directly compare the subspaces selected by the dimensionality reduction. We use a generalization of distance in the Grassmann space, which allows comparison of affine subspaces with unequal dimension [14]. The method augments the principle angles (defined in [22]) with enough additional angles (all equal to $\pi/2$) to increase the number of angles to the dimensionality of the larger space. An additional “direction vector” is also added to account for the affine offset.

We apply this generalization [14] to a new modification of the chordal distance, defined as

$$\theta_c = \arcsin \left(\left(\frac{1}{k} \sum_{i=1}^k \sin^2 \theta_i \right)^{1/2} \right), \quad (2)$$

for k principle angles $\{\theta_i\}_{i=1}^k$. The two modifications are 1) dividing by k , which allows the distance to be independent on the dimensionality of the spaces being compared, and 2) converting the distance measure back to an angle, which is more intuitive.

We then compare the subspaces obtained by PCA dimensionality reduction, for the same divisions of the data as used for Fig. 3. However, we now ignore any subsets of less than 50 HFOs, as these are deemed unreliable for computing the PCA in 33D. Note this figure still involves over 200 thousand time bin comparisons, the same 163 ictal versus interictal comparisons, and nearly 3,500 of each type of channel-channel comparisons.

Results for these comparisons are shown in Fig. 4. We observe that the PCA subspaces are quite consistent across different time bins, with the distributions for PCA subspaces all peaking less than 15° and not extending much past 20° .

4. PATHOLOGICAL VERSUS NORMAL HFOS

4.1. Dimensionality Reduction: Greedy Fisher LDA

While the subspaces obtained via PCA represent the variance well, they are not necessarily the optimal directions for separating the feature distributions of pHFOs and nHFOs. An alternate dimensionality reduction method is used: greedy Fisher's LDA. In this method, Fisher's LDA is applied, resulting in a single basis direction. The projection of the data in this direction is then subtracted from the data, and the process is repeated.

Recall that Fisher's LDA uses the sum of the covariance of each group, rather than the covariance of the pooled groups [23]. Letting the mean and covariance of the two groups be denoted μ_A , μ_B and S_A , S_B , respectively, the specific direction is given by

$$w \propto (\sum_A + \sum_B)^{-1} (\mu_A - \mu_B). \quad (3)$$

Note, the rank of the sum of the covariances will be reduced by one in each step. Thus, to invert the matrix, the eigenvalue method is used, with the inverse of the eigenvalues corresponding to the removed projections being set to zero.

We compute this basis, comparing ictal versus interictal times, for the 13 patients with recorded seizures. We select the number of basis vectors equal to the mean intrinsic dimension. We again compare the basis vectors using Eq. 2, with the results shown in Fig. 4. The LDA subspaces vary much more across channels than the PCA subspaces, with the θ_C distribution for LDA being almost fully localized between 30° and 50° . Note that 45° implies that the subspaces overlap by half. Thus, the Fisher LDA subspaces show that some differences in ictal versus interictal HFO features are consistent between channels, while other differences are not conserved.

4.2. Bayes Error Estimates of pHFOs versus nHFOs

Next we quantify how distinct the feature distributions of pHFOs and nHFOs are, per channel. This quantification serves as a guide for future work. We utilize the Henze-Penrose divergence (HPD) [16, 17] to compute bounds on the Bayes Error. Note that small upper bounds imply highly separable classes, whereas upper bounds near 0.5 imply inseparable classes. We estimate the HPD bound using the nonparametric ensemble estimator derived in [24, 25], which achieves the parametric convergence rate.

The left panel of Fig. 5 shows the Bayes Error bound estimates for the 163 channels (eight patients) with at least 50 HFOs in each of the ictal and interictal states. Channels are separated between SOZ (27 channels) and non-SOZ (124 channels) for patients with ILEA Class I (the best) surgery outcome, and an "other" category (12 channels), including channels from patients with either worse surgery outcomes, no surgery, or missing meta-data.

Patients with Both SOZ and non-SOZ channels in ILAE class I patients (best surgery outcome) have the Bayes Error rate bound to relatively small values. However, channels in other patients have a more diffuse distribution of upper bounds, with the most probable upper bound about 0.25. It is expected that ictal HFOs are almost entirely pHFOs (on any channel), that interictal HFOs on non-SOZ channels are predominately nHFOs, and that interictal HFOs on SOZ channels are a mixture of both pHFOs and nHFOs. Thus, it is expected that the Bayes Error bounds would be lower for non-SOZ channels. Overall, these bounds suggest that there is sufficient separation between pHFO and nHFO feature distributions to allow classification of pHFOs and nHFOs in most channels.

The right panel of Fig. 5 displays the lower bound versus a linear error estimate. The linear estimate was computed with 10-fold cross validation in the Fisher LDA space by using a “box” classification boundary, with the threshold in each dimension being the value for which the receiver operator curve has largest transverse distance from the diagonal. Linear regression was also performed, resulting in an offset of 0.06 (0.04–0.08 at 95% C.L.) and a slope of 1.05 (0.82–1.28 at 95% C.L.). Thus, in aggregate this “box” classifier is already relatively close to the bound on the Bayes Error, though many individual channels are still quite far from the bound.

5. CLUSTERING CHANNELS BASED ON SUBSPACES

Given the observed variations in greedy Fisher LDA subspaces across channels, we seek to compare the natural clustering of these subspaces with known groupings of channels. The most relevant groupings of channels are the groups based on the physical configuration of the recording electrodes (several grids or strips of electrodes are implanted for each patient), as well as the clinically determined SOZ and resected volume for ILAE Class I patients.

The unsupervised clustering is obtained by first converting the matrix of modified chordal distances (Eq. 2) between channels per each patient to a metric using the method of dual-rooted-trees followed by spectral clustering [26]. This method adapts to the natural geometry of the data and is competitive with other algorithms [26]. We select either two groups (as the SOZ and resected volume clustering is binary) or the number of groups equal to the number of strips and grids. The unsupervised clustering results are compared with these labels using the adjusted Rand index (ARI) [27]. Unfortunately, the requirement that there be at least 50 ictal HFOs reduces the number of channels per patient significantly, resulting in only a few (4–5) patients having enough channels to cluster. However, the ARI never exceeded 0.02 for any of these comparisons in any patients, suggesting that the primary distinction between channels is not the pathology or the grid/strip placing, but other effects.

6. DISCUSSION AND CONCLUSION

Overall, we observe that the HFO features tend to cluster on linear manifolds. Both the subspaces of these manifolds and the local intrinsic dimension tend to be consistent within interictal periods, but may change between interictal and ictal periods. We especially note significant differences in both the intrinsic dimension and feature manifolds between different channels within the same patient. Thus, dimensionality reduction and feature

analysis must account for variations between channels. The dominant cause of this inter-channel variation does not appear to be tissue pathology or grid/strip groups in the recording. We also observe that pHFOs and nHFOs are indeed distinct on a large number of channels, suggesting a strong potential for classifying individual HFOs.

This analysis also demonstrates methods applicable to other discrete events, including using the θ_I statistic for comparing local, intrinsic dimension for collections of events, an affine Grassmann distance θ_C for comparing consistency of subspaces, and estimating bounds on the Bayes Error to assess the feasibility of low-error classification. Future work will extend this analysis to a larger patient population, and classify HFOs and/or recording channels based on HFO features.

References

1. Park SC, Lee SK, Che H, Chung CK. Ictal high-gamma oscillation (60–99 Hz) in intracranial electroencephalography and postoperative seizure outcome in neocortical epilepsy. *Clin Neurophysiol.* 2012; 123(6):1100–10. [PubMed: 22391040]
2. Cho JR, Koo DL, Joo EY, Seo DW, Hong SC, Jiruska P, Hong SB. Resection of individually identified high-rate high-frequency oscillations region is associated with favorable outcome in neocortical epilepsy. *Epilepsia.* 2014; 55:1872–83. [PubMed: 25266626]
3. Haegelen C, Perucca P, Chillon CE, Andrade-Valena L, Zelman R, Jacobs J, Collins DL, Dubeau F, Olivier A, Gotman J. High-frequency oscillations, extent of surgical resection, and surgical outcome in drug-resistant focal epilepsy. *Epilepsia.* 2013; 54:848–57. [PubMed: 23294353]
4. Kerber K, Dimpelmann M, Schelter B, Le Van P, Korinthenberg R, Schulze-Bonhage A, Jacobs J. Differentiation of specific ripple patterns helps to identify epileptogenic areas for surgical procedures. *Clin Neurophysiol.* 2014; 125:1339–1345. [PubMed: 24368032]
5. Gliske S, Irwin Z, Chestek C, Stacey WC. Automated identification of seizure onset zone with a universal detector of high frequency oscillations. *Clin Neurophysiol.* 2015 INSERT.
6. Worrell GA, Jerbi K, Kobayashi K, Lina JM, Zelman R, Le Van Quyen M. Recording and analysis techniques for high-frequency oscillations. *Prog Neurobiol.* 2012; 98:265–278. [PubMed: 22420981]
7. Blanco JA, Stead M, Krieger A, Viventi J, Marsh WR, Lee KH, Worrell GA, Litt B. Unsupervised classification of high-frequency oscillations in human neocortical epilepsy and control patients. *J Neurophysiol.* 2010; 104:2900–2912. [PubMed: 20810694]
8. Blanco JA, Stead M, Krieger A, Stacey W, Maus D, Marsh E, Viventi J, Lee KH, Marsh R, Litt B, Worrell GA. Data mining neocortical high-frequency oscillations in epilepsy and controls. *Brain.* 2011; 134:2948–2959. [PubMed: 21903727]
9. Pearce A, Wulsin D, Blanco JA, Krieger A, Litt B, Stacey WC. Temporal changes of neocortical high-frequency oscillations in epilepsy. *J Neurophysiol.* 2013; 110:1167–1179. [PubMed: 23761699]
10. Matsumoto A, Brinkmann BH, Stead SM, Matsumoto J, Kuciewicz MT, Marsh WR, Meyer F, Worrell G. Pathological and physiological high-frequency oscillations in focal human epilepsy. *J Neurophysiol.* 2013; 110:1958–64. [PubMed: 23926038]
11. Carter KM, Raich R, Hero AO. On local intrinsic dimension estimation and its applications. *Signal Processing, IEEE Transactions on.* 2010; 58(2):650–663.
12. Costa, JA.; Hero, AO, III. *Statistics and Analysis of Shapes.* Springer; 2006. Determining intrinsic dimension and entropy of high-dimensional shape spaces; p. 231–252.
13. Moon KR, Li JJ, Delouille V, De Visscher R, Watson F, Hero AO III. Image patch analysis of sunspots and active regions. i. intrinsic dimension and correlation analysis. 2015 arXiv preprint arXiv:1503.04127.
14. Ye K, Lim L-H. Distance between subspaces of different dimensions. 2014 Preprint. arXiv: 1407.0900.

15. Wang J, Xu Y, Zhang D, You J. An efficient method for computing orthogonal discriminant vectors. *Neurocomputing*. 2010; 73:2168–76.
16. Moon, KR.; Delouille, V.; Hero, AO, III. Meta learning of bounds on the bayes classifier error. *IEEE Signal Processing and Signal Processing Education Workshop*; IEEE; 2015.
17. Berisha, V.; Wisler, A.; Hero, A., III; Spanias, A. Empirically estimable classification bounds based on a nonparametric divergence measure. *Signal Processing, IEEE Transactions on*; 2015;
18. Wagenaar JB, Worrell GA, Ives Z, Matthias D, Litt B, Schulze-Bonhage A. Collaborating and sharing data in epilepsy research. *J Clin Neurophysiol*. 2015; 32:235–9. [PubMed: 26035676]
19. Kaiser, JF. On a simple algorithm to calculate the energy of a signal. *IEEE Int. Conf. Acoustic Speech Signal Process*; IEEE; 1990.
20. Ochiai A. Zoogeographical studies on the soleoid fishes found japan and its neighboring regions. ii. *Bull Jap Soc Sci Fish*. 1957; 22(9):526–530.
21. Barkman JJ. Phytosociology and ecology of cryptogamic epiphytes, including a taxonomic survey and description of their vegetation units in europe. Assen Van Gorcum. 1958
22. Hotelling H. Relations between two sets of variates. *Biometrika*. 1936; 28:321–377.
23. Fisher RA. The use of multiple measurements in taxonomical problems. *Annals of Eugenics*. 1936; 7(2):179–188.
24. Moon, KR.; Hero, AO. Ensemble estimation of multivariate f-divergence. *Information Theory (ISIT), 2014 IEEE International Symposium on*; IEEE; 2014. p. 356-360.
25. Moon K, Hero A. Multivariate f-divergence estimation with confidence. *Advances in Neural Information Processing Systems*. 2014:2420–2428.
26. Galluccio L, Michel O, Comon P, Kliger M, Hero AO III. Clustering with a new distance measure based on a dual-rooted tree. *Information Sciences*. 2013; 251:96–113.
27. Rand WM. Objective criteria for the evaluation of clustering methods. *Journal of the American Statistical Association*. 1971; 66(336):846–850.

Fig. 1.

Diagram relating the recorded data with the implanted electrodes. A 5×7 grid of electrodes placed over a region of cortex. Each channel produces a separate time series of data, with some channels being identified by clinicians as seizure onset zone (SOZ). HFO detections are marked by solid magenta lines under the EEG trace.

Fig. 2.

Example HFO detections within 45 min of one channel of intracranial EEG data. A seizure occurs at about 35 minutes. HFO detections (72 interictal and preictal, 32 ictal) are shown as small yellow dots. Two HFOs are also shown using a *much* smaller scale.

Fig. 3. Comparisons of the consistency of the intrinsic dimension (θ_I from Eq. 1) for four comparisons, as described in the text. Note, 0° implies no difference and 90° implies maximal difference between the two collections of intrinsic dimension being compared.

Fig. 4. Comparison of the subspaces from the PCA and greedy Fisher LDA dimensionality reduction, using θ_C (Eq. 2), for the same types of comparisons as Fig. 3. Again, 0° implies maximally similar and 90° implies maximally different.

Fig. 5. Distribution of upper Bayes Error bound based on Henze-Penrose divergences (HPD) between ictal and interictal HFOs per channel. Left panel: upper bounds stratified by channel type. Bounds near zero imply high distinction between pHFOs and nHFOs, and bounds near 0.5 imply no distinction. Right panel: comparison between lower bound estimate and a linear error estimate (see the text), including the linear regression fit.