
Forecastable Component Analysis

Georg M. Goerg

GMG@STAT.CMU.EDU

Carnegie Mellon University, Department of Statistics, Pittsburgh, PA 15213

Abstract

I introduce **Forecastable Component Analysis (ForeCA)**, a novel dimension reduction technique for temporally dependent signals. Based on a new *forecastability* measure, ForeCA finds an optimal transformation to separate a multivariate time series into a forecastable and an orthogonal white noise space. I present a converging algorithm with a fast eigenvector solution. Applications to financial and macro-economic time series show that ForeCA can successfully discover informative structure, which can be used for forecasting as well as classification.

The R package `ForeCA` accompanies this work and is publicly available on [CRAN](#).

1. Introduction

With the rise of high-dimensional datasets it has become important to perform dimension reduction (DR) to a lower dimensional representation of the data. For simplicity we consider linear transformations $\mathbf{W} \in \mathbb{R}^{k \times n}$, which map an n -dimensional \mathbf{X} to a $k \leq n$ dimensional $\mathbf{S} = \mathbf{W}\mathbf{X}$. Typically, the transformed data should be somewhat “interesting”; there is no point in transforming \mathbf{X} to an arbitrary \mathbf{S} that is less useful, meaningful, etc. Let $\iota(\mathbf{S})$ measure “interestingness” of \mathbf{S} . DR can then be set up as an optimization problem

$$\hat{\mathbf{w}}_j = \arg \max_{\mathbf{w} \in \mathbb{R}^{n \times 1}} \iota(\mathbf{w}^\top \mathbf{X}), \quad j = 1, \dots, k, \quad (1)$$

$$\text{subject to } \mathbf{w}_j^\top \mathbf{X} \perp \{\mathbf{w}_1^\top \mathbf{X}, \dots, \mathbf{w}_{j-1}^\top \mathbf{X}\}, \quad (2)$$

where (2) is a common DR constraint, which makes $\mathbf{S}_j = \mathbf{w}_j^\top \mathbf{X}$ orthogonal (uncorrelated) to previously obtained signals.

For example, principal component analysis (PCA) keeps large variance signals (Jolliffe, 2002) – $\iota(X) =$

$\mathbb{E}(X - \mathbb{E}X)^2$ in (1); independent component analysis (ICA) recovers statistically independent signals (Hyvärinen and Oja, 2000); slow feature analysis (SFA) (Wiskott and Sejnowski, 2002) finds “slow” signals and is equivalent to maximizing the lag 1 autocorrelation coefficient.

DR techniques are often applied to multivariate time series \mathbf{X}_t , hoping that forecasting on the lower-dimensional space \mathbf{S}_t is more accurate, simpler, more efficient, etc. Standard DR techniques such as PCA or ICA, however, do not explicitly address forecastability of the sources. For example, just because a signal has high variance does not mean it is easy to forecast.

Thus let’s define interesting as being predictable. Forecasting is not only good for its own sake (finance, economics), but even when future values are not immediately interesting, signals that do have predictive power exhibit non-trivial structure by definition – and are thus easier to interpret. For example, the time series in Fig. 1 are ordered from least (S&P500 daily returns) to most forecastable (monthly temperature in Nottingham) according to the ForeCA forecastability measure $\Omega(x_t)$ I propose in Definition 3.1 below. And indeed moving from left to right they exhibit more structure.

The main contributions of this work are i) a model-free, comparable measure of forecastability for (stationary) time series (Section 3), ii) a novel data-driven DR technique, ForeCA, that finds forecastable signals, iii) an iterative algorithm that provably converges to (local) optima using fast eigenvector solutions (Section 4), and iv) applications showing that ForeCA outperforms traditional DR techniques in finding low-dimensional, forecastable subspaces, and that it can also be used for time series classification (Section 5). Related work will be reviewed in Section 6.

All computations and simulations were done in R (R Development Core Team, 2010).

2. Time Series Preliminaries

Let y_t be a univariate, second-order stationary time series with mean $\mathbb{E}y_t = \mu_y < \infty$, variance $\mathbb{V}y_t = \sigma_y^2$,

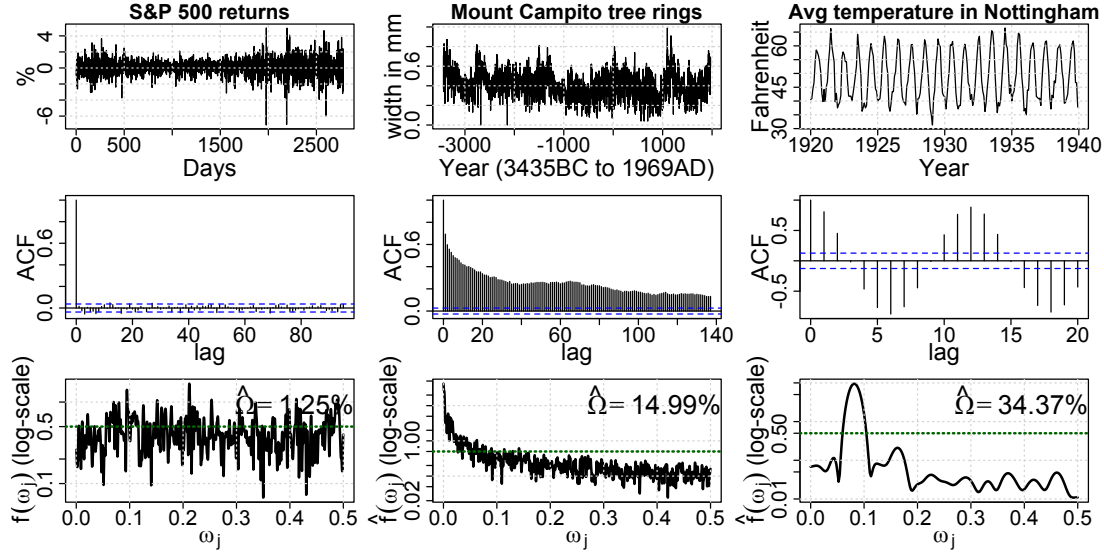


Figure 1. Observations (top); sample ACF $\hat{\rho}(k)$ (middle); smoothed WOSA spectral density estimate (bottom). From left to right: i) S&P 500 daily returns; ii) Mount Campito tree ring series; iii) monthly mean temperatures in Nottingham. Data publicly available in R packages: SP500 in MASS; camp in tseries; nottem in datasets.

and autocovariance function (ACVF)

$$\gamma_y(k) = \mathbb{E}(y_t - \mu_y)(y_{t-k} - \mu_y), \quad k \in \mathbb{Z}. \quad (3)$$

The ACVF for univariate processes is symmetric in k , $\gamma_y(k) = \gamma_y(-k)$. Let $\rho(k) = \gamma(k)/\gamma(0)$ be the autocorrelation function (ACF). A large $\rho(k)$ means that the process k time steps ago is highly correlated with the present y_t . The sample ACFs $\hat{\rho}(k)$ in Fig. 1 show that, e.g., S&P 500 daily returns are uncorrelated with their own past (stock market efficiency); yearly tree ring growth is highly correlated over time with significant lags even for $k \geq 100$ years; and intuitively temperature in month t is highly correlated with the temperature $k = 6$ (cold \leftrightarrow warm) and $k = 12$ (cold \rightarrow cold; warm \rightarrow warm) months ago (or in the future).

The building block of time series models is *white noise* ε_t , which has zero mean, finite variance, and is uncorrelated over time: $\varepsilon_t \sim WN(0, \sigma_\varepsilon^2)$ iff¹ i) $\mathbb{E}\varepsilon_t = 0$, ii) $\mathbb{V}\varepsilon_t = \gamma_\varepsilon(0) = \sigma_\varepsilon^2$, and iii) $\gamma_\varepsilon(k) = 0$ if $k \neq 0$. Only if ε_t is a Gaussian process, then it is also independent.

For multivariate second-order stationary \mathbf{X}_t with mean² $\boldsymbol{\mu} \in \mathbb{R}^n$ and covariance matrix $\Sigma_{\mathbf{X}}$ the ACVF

$$\mathbb{R}^{n \times n} \ni \boldsymbol{\Gamma}_{\mathbf{X}}(k) = \mathbb{E}(\mathbf{X}_t - \boldsymbol{\mu})(\mathbf{X}_{t-k} - \boldsymbol{\mu})^\top, \quad (4)$$

is a matrix-valued function of $k \in \mathbb{Z}$. In particular, $\boldsymbol{\Gamma}_{\mathbf{X}}(0) = \Sigma_{\mathbf{X}}$. The diagonal of $\boldsymbol{\Gamma}_{\mathbf{X}}(k)$ contains the ACVF of each $\mathbf{X}_i(t)$; the off-diagonal element

$\boldsymbol{\Gamma}_{\mathbf{X}}(k)_{(i,j)}$ is the cross-covariance between the i th and j th series at lag k :

$$\gamma_{ij}(k) = \mathbb{E}(\mathbf{X}_{i,t} - \mu_i)(\mathbf{X}_{j,t-k} - \mu_j) \in \mathbb{R}. \quad (5)$$

Contrary to $\gamma_y(k)$, $\boldsymbol{\Gamma}_{\mathbf{X}}(k)$ is not symmetric, but

$$\boldsymbol{\Gamma}_{\mathbf{X}}(k) = \boldsymbol{\Gamma}_{\mathbf{X}}(-k)^\top. \quad (6)$$

2.1. Spectrum and Spectral Density

The spectrum of a univariate stationary process can be defined as the Fourier transform of its ACVF,

$$S_y(\lambda) = \frac{1}{2\pi} \sum_{j=-\infty}^{\infty} \gamma_y(j) e^{ij\lambda}, \quad \lambda \in [-\pi, \pi], \quad (7)$$

where $i = \sqrt{-1}$ is the imaginary unit. Since $\gamma_y(k)$ is symmetric, the spectrum is a real-valued, non-negative function, $S_y : [-\pi, \pi] \rightarrow \mathbb{R}^+$. For white noise ε_t all $\gamma_\varepsilon(k) = 0$ if $k \neq 0$, thus $S_\varepsilon(\lambda) = \frac{\sigma_\varepsilon^2}{2\pi}$ is constant for all $\lambda \in [-\pi, \pi]$. When $\gamma(k) > 0$ for $k \neq 0$ the spectrum has peaks at the corresponding frequencies. For example, the spectral density of monthly temperature series (right in Fig. 1) has large peaks at $\lambda \approx \pi/6$ and $\pi/12$, which represent the half- and one-year cycle.³

Vice versa, the ACVF can be recovered from the spec-

¹Iff will be used as an abbreviation for *if and only if*.

²Without loss of generality (WLOG) assume $\boldsymbol{\mu} = \mathbf{0}$.

³Frequencies λ are often scaled by π , $\tilde{\lambda} = \lambda/\pi$. This does not change results qualitatively, but simplifies interpretation since the corresponding cycle length equals $\tilde{\lambda}^{-1}$.

trum using the inverse Fourier transform

$$\gamma_y(k) = \int_{-\pi}^{\pi} S_y(\lambda) e^{-ik\lambda} d\lambda, \quad k \in \mathbb{Z}. \quad (8)$$

In particular, $\int_{-\pi}^{\pi} S_y(\lambda) d\lambda = \sigma_y^2$ for $k = 0$. Let

$$f_y(\lambda) = \frac{S_y(\lambda)}{\sigma_y^2} = \frac{1}{2\pi} \sum_{j=-\infty}^{\infty} \rho_y(j) e^{ij\lambda}, \quad (9)$$

be the *spectral density* of y_t . As $f_y(\lambda) \geq 0$ and $\int_{-\pi}^{\pi} f_y(\lambda) d\lambda = 1$, the spectral density can be interpreted as a probability density function (pdf) of an (unobserved) random variable (RV) Λ that “lives” on the unit circle. For white noise $f_\varepsilon(\lambda) = \frac{1}{2\pi}$, which represents the uniform distribution $U(-\pi, \pi)$.

Remark 2.1 (Spectrum and spectral density). *In the time series literature “spectrum” and “spectral density” are often used interchangeably. Here I reserve “spectral density” for $f_y(\lambda)$ in (9), as it integrates to one such as standard probability density functions.*

3. Measuring Forecastability

Forecasting is inherently tied to the time domain. Yet, since Eqs. (7) & (8) provide a one-to-one mapping between the time and frequency domain, we can use frequency domain properties to measure forecastability.

The intuition for the proposed measure of forecastability is as follows. Consider

$$\begin{aligned} y_t &= \sqrt{2} \cos(2\pi\mathcal{Y}t + \theta), \\ \theta &\sim U(-\pi, \pi), \quad \mathcal{Y} \sim p_y(y) \text{ independent of } \theta. \end{aligned} \quad (10)$$

One can show that $S_y(\lambda) = p_y(\lambda)$ (Gibson, 1994).

If we have to predict the future of y_t , then uncertainty about y_{t+h} , $h > 0$, is only manifested in uncertainty about \mathcal{Y} , since $\cos(2\pi\mathcal{Y}t + \theta)$ is a deterministic function of t : less uncertainty about \mathcal{Y} means less uncertainty about y_{t+h} . We can measure this uncertainty using the Shannon entropy of $p_y(y)$ (Shannon, 1948). It is thus natural to measure uncertainty about the future as (differential) entropy of $f_y(\lambda)$,

$$H_{s,a}(y_t) := - \int_{-\pi}^{\pi} f_y(\lambda) \log_a f_y(\lambda) d\lambda, \quad (11)$$

where $a > 0$ is the logarithm base.

On a finite support $[b, c]$ the maximum entropy occurs for the uniform distribution $U(b, c)$; thus a flat spectrum should indicate the least predictable sequence. And indeed, a flat spectrum corresponds to white

noise, which is unpredictable by definition (using linear predictors). Consequently, for any stationary y_t

$$\begin{aligned} H_{s,a}(y_t) &\leq H_{s,a}(\text{white noise}) \\ &= - \int_{-\pi}^{\pi} \frac{1}{2\pi} \log_a \frac{1}{2\pi} d\lambda = \log_a 2\pi, \end{aligned}$$

with equality iff y_t is white noise.

Definition 3.1 (Forecastability of a stationary process). *For a second-order stationary process y_t , let*

$$\begin{aligned} \Omega : y_t &\mapsto [0, \infty], \\ \Omega(y_t) &= 1 - \frac{H_{s,a}(y_t)}{\log_a(2\pi)} = 1 - H_{s,2\pi}(y_t), \end{aligned} \quad (12)$$

be the forecastability of y_t .

Contrary to other measures in the signal processing and time series literature, $\Omega(y_t)$ does not require actual forecasts, but is a characteristic of the process y_t . It is therefore not biased to a particular – perhaps sub-optimal – model, forecast horizon, or loss function; as used in e.g., Box and Tiao (1977); Stone (2001).

Properties 3.2. $\Omega(y_t)$ satisfies:

a) $\Omega(y_t) = 0$ iff y_t is white noise.

b) invariant to scaling and shifting:

$$\Omega(ay_t + b) = \Omega(y_t) \text{ for } a, b \in \mathbb{R}, a \neq 0.$$

c) max sub-additivity for uncorrelated processes:

$$\Omega(\alpha x_t + \sqrt{1 - \alpha^2} y_t) \leq \max\{\Omega(x_t), \Omega(y_t)\}, \quad (13)$$

if $\mathbb{E}x_t y_s = 0$ for all $s, t \in \mathbb{Z}$; equality iff $\alpha \in \{0, 1\}$.

The three series in Fig. 1 are ordered (left to right) by increasing forecastability and indeed larger $\hat{\Omega}$ correspond to intuitively more predictable real-world events: stock returns are in general not predictable; average monthly temperature is.

We can thus use (12) to guide the search for optimal \mathbf{w} that make $y_t = \mathbf{w}^\top \mathbf{X}_t$ as forecastable as possible.

3.1. Plug-in Estimator for Ω

To estimate $\Omega(y_t)$, we first estimate $S_y(\lambda)$, normalize it, and then plug it in (11).

An unbiased estimator of $S_y(\lambda)$ is the periodogram

$$I_{T, \mathbf{y}_1^T}(\omega_j) = \left| \frac{1}{\sqrt{T}} \sum_{t=0}^{T-1} y_t e^{-2\pi i \omega_j t} \right|^2, \quad (14)$$

where $\omega_j = j/T$, $j = 0, 1, \dots, T-1$ are the (scaled) Fourier frequencies, and $\mathbf{y}_1^T = \{y_1, \dots, y_T\}$ is a sample of y_t . It is well known that (14) is not a good estimate (e.g., periodograms are not consistent). In the numerical examples we therefore use weighted overlapping segment averaging (WOSA) (Nuttall and Carter, 1982) $\widehat{S}_y(\omega_j)$ from the R package `sapa: SDF(y, 'wosa')`.

The bottom row of Figure 1 shows the normalized $\widehat{f}_{j,y} = \frac{\widehat{S}_y(\omega_j)}{\sum_{j=0}^{T-1} \widehat{S}_y(\omega_j)}$ along with the plug-in estimate

$$\widehat{\Omega}(\mathbf{y}_1^T) = 1 + \sum_{j=0}^{T-1} \widehat{f}_{j,y} \cdot \log_{a=T}(\widehat{f}_{j,y}). \quad (15)$$

Remark 3.3. Typically, to estimate $\mathbb{E}g(X)$ for $X \sim p(x)$ (here: $g(X) = \log p(X)$) the sample average is solely over $g(x_j)$ without multiplicative $p(x_j)$ terms. This however assumes that each x_j is sampled from $p(x)$ (and thus $\frac{1}{n} \sum_{i=1}^n g(x_i) \rightarrow \mathbb{E}_p g(X) = \int g(x)p(x)dx$ by the strong law of large numbers). While this is true in a standard sampling framework, here the “data” are the Fourier frequencies ω_j and the fast Fourier transform (FFT) samples them uniformly (and deterministically) from $[-\pi, \pi]$ and not according to the “true” spectral density $f(\lambda)$.⁴

Eq. (15) can be improved by a better spectral density (Fryzlewicz, Nason, and von Sachs, 2008; Lees and Park, 1995; Trobs and Heinzl, 2006) and entropy estimation (Paninski, 2003). Future research can also address direct estimation of (11) – as is common for classic entropy estimates (Sricharan, Raich, and Hero, 2011; Stowell and Plumbley, 2009). However, since neither spectrum nor entropy estimation are the primary focus of this work, we use standard estimators for $S_y(\lambda)$ and then the plug-in estimator of (15).

It must be noted though that $\widehat{\Omega}(\mathbf{y}_1^T)$ in (15) is based on discrete rather than differential entropy. It still has the intuitive property that white noise has zero estimated forecastability, but now $\widehat{\Omega}(\mathbf{y}_1^T) \in [0, 1]$; $\widehat{\Omega}(\mathbf{y}_1^T) = 1$ iff the sample is a perfect sinusoid. Applications show that (15) yields reasonable estimates and we do not expect the results to change qualitatively for other estimators. We leave differential entropy estimates of Ω to future work.

Notice that $\Omega(y_t)$ relies on Gaussianity as only then $f_y(\lambda)$ captures all the temporal dependence structure of y_t . While time series are often non-Gaussian, $\Omega(\cdot)$ is a computationally and algebraically manageable forecastability measure – similarly to the importance of

variance in PCA for iid data, even though they are rarely Gaussian.

4. ForeCA: Maximizing Forecastability

Recall from Eq. (1) that we want to find a linear combination of a multivariate \mathbf{X}_t that makes $y_t = \mathbf{w}^\top \mathbf{X}_t$ as forecastable as possible. Based on the forecastability measure in Section 3, we can now formally define the ForeCA optimization problem:

$$\max_{\mathbf{w}} \Omega(\mathbf{w}^\top \mathbf{X}_t) = \max_{\mathbf{w}} \left(1 + \frac{\int_{-\pi}^{\pi} f_y(\lambda) \log_a f_y(\lambda) d\lambda}{\log_a(2\pi)} \right), \quad (16)$$

$$\text{subject to } \mathbf{w}^\top \Sigma_{\mathbf{X}} \mathbf{w} = 1, \quad (17)$$

where (17) must hold since (11) uses the spectral density of y_t , i.e. we need $\mathbb{V}y_t = \mathbf{w}^\top \Sigma_{\mathbf{X}} \mathbf{w} = 1$.

Property 3.2c seems to let (16) only have a trivial boundary solution. However, it is intuitively clear that combining uncorrelated series makes forecasting (in general) more difficult, e.g., signal + noise. But if $\mathbb{E}x_t y_s \neq 0$ for some $s, t \in \mathbb{Z}$ then combining them can make it simpler: for some $\alpha \in (0, 1)$ it holds $\Omega(\alpha x_t + \sqrt{1-\alpha^2} y_t) > \max\{\Omega(x_t), \Omega(y_t)\}$.

To optimize the right hand side of (16) we need to evaluate $f_y(\lambda) = f_{\mathbf{w}^\top \mathbf{X}_t}(\lambda)$ for various \mathbf{w} and do this efficiently. We now show how to obtain $f_y(\lambda)$ by simple matrix-vector multiplication from $f_{\mathbf{X}}(\lambda)$.

4.1. Spectrum of Multivariate Time Series and Their Linear Combinations

For multivariate \mathbf{X}_t the spectrum equals

$$S_{\mathbf{X}}(\lambda) = \frac{1}{2\pi} \sum_{k=-\infty}^{\infty} \Gamma_{\mathbf{X}}(k) e^{2\pi i k \lambda}, \quad \lambda \in [-\pi, \pi]. \quad (18)$$

Contrary to the univariate case, (18) is in general complex-valued. Yet, since $\Gamma_{\mathbf{X}}(k) = \Gamma_{\mathbf{X}}(-k)^\top$, $S_{\mathbf{X}}(\lambda) \in \mathbb{C}^{n \times n}$ is Hermitian for every λ , $S_{\mathbf{X}}(\lambda) = \overline{S_{\mathbf{X}}(\lambda)^\top}$, where $\bar{z} = a - ib$ is the complex conjugate of $z = a + ib \in \mathbb{C}$ (Brockwell and Davis, 1991, p. 436).

For dimension reduction we consider linear combinations $y_t = \mathbf{w}^\top \mathbf{X}_t$, $\mathbf{w} \in \mathbb{R}^n$. By assumption $\mathbb{E}y_t = \mathbf{w}^\top \mathbb{E}\mathbf{X}_t = 0$ and $\gamma_y(k) = \mathbb{E}y_t y_{t-k} = \mathbf{w}^\top \Gamma_{\mathbf{X}}(k) \mathbf{w}$. In particular, $\gamma_y(0) = \sigma_y^2 = \mathbf{w}^\top \Sigma_{\mathbf{X}} \mathbf{w}$. The spectrum of $\mathbf{w}^\top \mathbf{X}_t$ can be quickly computed via $S_y(\lambda) = \mathbf{w}^\top S_{\mathbf{X}}(\lambda) \mathbf{w}$ and consequently

$$f_y(\lambda) = \frac{\mathbf{w}^\top S_{\mathbf{X}}(\lambda) \mathbf{w}}{\mathbf{w}^\top \Sigma_{\mathbf{X}} \mathbf{w}}, \quad \lambda \in [-\pi, \pi]. \quad (19)$$

⁴Advances in “compressed sensing” (Jacques and Vandergheynst, 2010) might improve estimates; see also “non-uniform FFT” (Fessler and Sutton, 2003).

Since $f_y(\lambda) \geq 0$ for every y_t , $\mathbf{w}^\top S_{\mathbf{X}}(\lambda)\mathbf{w} \geq 0$ for all $\mathbf{w} \in \mathbb{R}^n$; thus $S_{\mathbf{X}}(\lambda)$ is positive semi-definite.

4.2. Solving the Optimization Problem

Since Ω is invariant to shift and scale (Property 3.2b), we shall not only assume zero mean, but also contemporaneously uncorrelated observed signals with unit variance in each component. WLOG consider $\mathbf{U}_t = \Sigma_{\mathbf{X}}^{-1/2}\mathbf{X}_t$; thus $\mathbb{E}\mathbf{U}_t\mathbf{U}_t^\top = \mathbf{I}_n$. Given $\widehat{\mathbf{W}}_U$ for \mathbf{U}_t , the transformation for \mathbf{X}_t becomes $\widehat{\mathbf{W}}_X = \widehat{\mathbf{W}}_U\widehat{\Sigma}_X^{-1/2}$. Problem (16) is then equivalent to

$$\mathbf{w}^* = \arg \min_{\mathbf{w}, \|\mathbf{w}\|_2=1} h(\mathbf{w}) \quad (20)$$

where

$$h(\mathbf{w}) = - \int_{-\pi}^{\pi} \mathbf{w}^\top S_U(\lambda)\mathbf{w} \cdot \ell(\mathbf{w}; \lambda) d\lambda, \quad (21)$$

is the spectral entropy (Eq. (11)) of $\mathbf{w}^\top \mathbf{X}_t$ as a function of \mathbf{w} . We use $\ell(\mathbf{w}; \lambda) := \log \mathbf{w}^\top S_U(\lambda)\mathbf{w} = \log f_{\mathbf{w}^\top \mathbf{U}}(\lambda)$ for better readability.

In practice we approximate (21) with $\widehat{S}_U(\omega_j) \in \mathbb{C}^{n \times n}$ and thus obtain⁵

$$\mathbf{w}^* = \arg \min_{\mathbf{w}, \|\mathbf{w}\|_2=1} \widehat{h}_T(\mathbf{w}). \quad (22)$$

Here

$$\widehat{h}_T(\mathbf{w}) = -\frac{1}{T} \sum_{j=1}^{T-1} \left(\mathbf{w}^\top \widehat{S}_U(\omega_j)\mathbf{w} \right) \cdot \widehat{\ell}(\mathbf{w}; \omega_j) \quad (23)$$

is the discretized version of (20), where $\widehat{\ell}(\mathbf{w}; \omega_j) = \log \mathbf{w}^\top \widehat{S}_U(\omega_j)\mathbf{w}$. Notice that $\widehat{S}_U(\omega_j) \in \mathbb{C}^{n \times n}$ varies with ω_j while $\mathbf{w} \in \mathbb{R}^n$ is fixed over all frequencies, which makes it difficult to obtain an analytic, closed-form solution. However, (22) can be solved iteratively borrowing ideas from the expectation maximization (EM) algorithm (Dempster, Laird, and Rubin, 1977).

4.2.1. A CONVERGENT EM-LIKE ALGORITHM

For every $\mathbf{w} \in \mathbb{R}^n$, $\|\mathbf{w}\|_2 = 1$, $h(\mathbf{w})$ has the form of a mixture model with weights $\widehat{\pi}(j | \mathbf{w}) := \mathbf{w}^\top \widehat{S}_U(\omega_j)\mathbf{w} \geq 0$ and ‘‘log-likelihood’’ $\widehat{\ell}(\mathbf{w}; \omega_j)$. Since $\int_{-\pi}^{\pi} f_{\mathbf{w}^\top \mathbf{U}}(\lambda) d\lambda = 1$, $\widehat{\pi}(j | \mathbf{w})$ is indeed a discrete probability distribution over $\{\omega_j | 0 = 1, \dots, T-1\}$.

Just as in an EM algorithm, the objective $h(\mathbf{w})$ can be optimized iteratively by first fixing $\mathbf{w} \leftarrow \mathbf{w}^{(i)}$ in

⁵We use ‘‘wosa’’ estimates (sapa R package). However, any other estimate of $S_U(\lambda)$ can be used.

$\widehat{\ell}(\mathbf{w}; \omega_j)$, and then minimizing the quadratic form

$$\mathbf{w}_{i+1} = \arg \min_{\mathbf{w}, \|\mathbf{w}\|_2=1} \mathbf{w}^\top \widehat{S}_U^{(i)}\mathbf{w}, \quad (24)$$

where $\widehat{S}_U^{(i)} = -\frac{1}{T} \sum_{j=0}^{T-1} \widehat{S}_U(\omega_j) \cdot \ell(\mathbf{w}_i; \omega_j)$.

Proposition 4.1. $\widehat{S}_U^{(i)}$ is positive semi-definite.

Thus (24) can be solved analytically by the last eigenvector of $\widehat{S}_U^{(i)}$ – automatically guaranteeing $\|\mathbf{w}\|_2 = 1$. The procedure iterates until $\|\mathbf{w}_{i+1} - \mathbf{w}_i\| < tol$ for some tolerance level tol . For initialization we sample \mathbf{w}_0 from an n -dimensional uniform hyper-cube, $U_n(-1, 1)$, and normalize to $\mathbf{w}_0 = \mathbf{w}_0 / \sqrt{\sum_{j=1}^n w_{j,0}^2}$.

Theorem 4.2 (Convergence). *The sequence $\{\mathbf{w}_i\}_{i \geq 0}$ obtained via (24) converges to a local minimum $\widehat{h}_T(\mathbf{w}^*) = \lambda_{\min}^* \geq 0$, where $\lim_{i \rightarrow \infty} \mathbf{w}_i = \mathbf{w}^*$ and λ_{\min}^* is the smallest eigenvalue of $\widehat{S}_U^{(*)}$.*

Corollary 4.3. *The transformed data $\mathbf{y}_1^{T, (*)} = \mathbf{w}^{(*)\top} \mathbf{X}_1^T$ satisfies*

$$\widehat{\Omega} \left(\mathbf{y}_1^{T, (*)} \right) = 1 - \lambda_{\min}^*. \quad (25)$$

Proof of Theorem 4.2. The entropy of a RV taking values in a finite alphabet $\{\omega_0, \dots, \omega_{T-1}\}$ is bounded: $0 \leq \widehat{h}_T(\mathbf{w}) \leq \log_a T$ for all $\mathbf{w} \in \mathbb{R}^n$. For convergence it remains to be shown that $\widehat{h}_T(\mathbf{w}_i) \geq \widehat{h}_T(\mathbf{w}_{i+1})$ with equality iff $\mathbf{w}_{i+1} = \mathbf{w}_i = \mathbf{w}^*$. First,

$$\begin{aligned} \widehat{h}_T(\mathbf{w}_i) &= -\frac{1}{T} \sum_{j=1}^{T-1} \mathbf{w}_i^\top S_U(\omega_j)\mathbf{w}_i \cdot \widehat{\ell}(\mathbf{w}_i; \omega_j) \\ &= \mathbf{w}_i^\top \widehat{S}_U^{(i)}\mathbf{w}_i \geq \mathbf{w}_{i+1}^\top \widehat{S}_U^{(i)}\mathbf{w}_{i+1} \end{aligned} \quad (26)$$

since \mathbf{w}_{i+1} is the last eigenvector of $\widehat{S}_U^{(i)}$. Second,

$$\begin{aligned} \mathbf{w}_{i+1}^\top \widehat{S}_U^{(i)}\mathbf{w}_{i+1} &= -\frac{1}{T} \sum_{j=1}^{T-1} \mathbf{w}_{i+1}^\top S_U(\omega_j)\mathbf{w}_{i+1} \cdot \widehat{\ell}(\mathbf{w}_i; \omega_j) \\ &\geq -\frac{1}{T} \sum_{j=1}^{T-1} \mathbf{w}_{i+1}^\top S_U(\omega_j)\mathbf{w}_{i+1} \cdot \widehat{\ell}(\mathbf{w}_{i+1}; \omega_j) \\ &= \widehat{h}_T(\mathbf{w}_{i+1}), \end{aligned} \quad (27)$$

where (27) holds as $\mathbb{E}_p - \log q = -\sum_{j=1}^n p_j \log q_j \geq -\sum_{j=1}^n p_j \log p_j = \mathbb{E}_p - \log p$ for any $q \neq p$. \square

To lower the chance of landing in local optima we repeat (24) for several random starting positions \mathbf{w}_0 and then select the best solution.

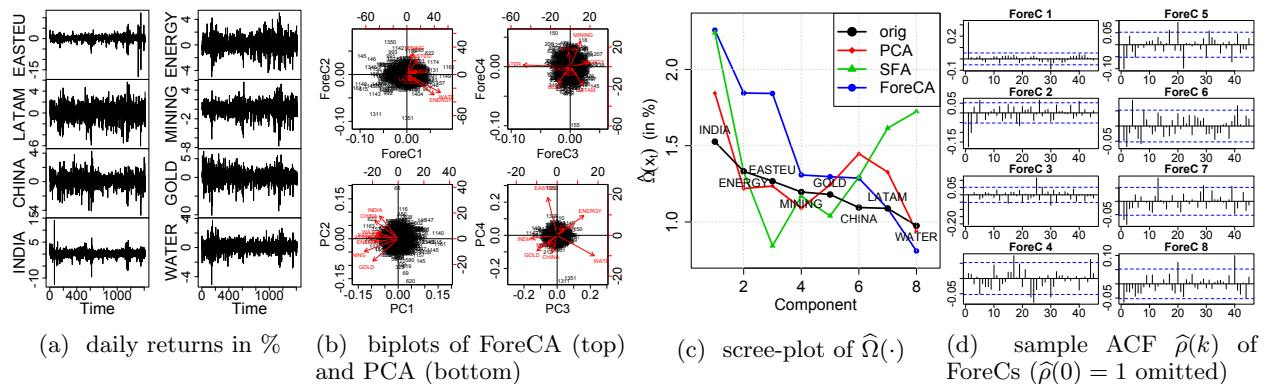


Figure 2. Equity fund returns analyzed with PCA, SFA, and ForeCA. (Dataset `equityFunds` in R package `fEcofin`.)

4.3. Obtaining a K -dimensional Subspace

To obtain all K loadings $\mathbf{W}_{1,\dots,K} = [\mathbf{w}_1, \dots, \mathbf{w}_K]$ that give uncorrelated series $y_{j,t}$, we iteratively (starting at $k = 1$) i) compute \mathbf{w}_k , ii) project \mathbf{U} onto the null space of $\mathbf{W}_{1,\dots,k} \rightarrow \mathbf{U}^{(k)} = \mathbf{W}_{1,\dots,k}^\perp \mathbf{U} \in \mathbb{R}^{K-k}$, iii) apply the EM-type algorithm on $\mathbf{U}^{(k)}$ to obtain $\tilde{\mathbf{w}}_{k+1}$, and finally iv) transform $\tilde{\mathbf{w}}_{k+1}$ back to loadings $\mathbf{w}^{(k)}$ of \mathbf{U} .

Doing this for $k = 1, \dots, K$ gives K loadings $\widehat{\mathbf{W}}_U$. Loadings for \mathbf{X}_t are given by $\widehat{\mathbf{W}}_X = \widehat{\mathbf{W}}_U \widehat{\Sigma}_X^{-1/2}$.

5. Applications

Here we demonstrate the usefulness of ForeCA to find informative, forecastable signals, but also as a tool for time series classification.

5.1. Improving Portfolio Forecasts

Figure 2a shows daily returns of eight equity funds from 2002/01/01 to 2007/05/31 ($T = 1413$). In the financial context finding forecastable series is an important goal by itself, not just for structure discovery. In particular, we can interpret a linear combination \mathbf{w} as a portfolio of stocks. The \mathbf{w}^* with the highest Ω gives the most forecastable portfolio.

Figure 2b shows a bi-plot for PCA and ForeCA for $(\mathbf{w}_1, \mathbf{w}_2)$ and $(\mathbf{w}_3, \mathbf{w}_4)$. As PC 1 weighs all funds almost equally, it represents the average market movement; the second component contrasts Gold & Mining with the rest and we can therefore label PC 2 as the “commodity” index. The third and fourth PC indicate energy/infrastructure and geographic regions.

However, even though PC 1 is also the most predictable PC, it has only a slightly larger $\hat{\Omega}$ than the most forecastable fund, India (Fig. 2c). On the other hand, combining Water (weight $w_{water,1} = 0.72$) with

Energy (0.58) is almost twice as forecastable as India (weights are from ForeC 1 in Fig. 2b). ForeC 2 also has high forecastability by selling Energy & Water (-0.53 & -0.47) and buying Mining & Eastern Europe (0.55 & 0.38). The third and fourth ForeCs seem to be hedging strategies (ForeC 3: Water vs. Energy; ForeC 4: Latin America & Gold vs. China & Mining).

As financial data only has very small autocorrelation – and usually at lag 1, if any –, SFA and ForeCA yield overall very similar results, except for a “wrong” ranking by SFA (Fig. 2c): SF 8 is the fastest feature (large, but negative lag 1 autocorrelation), yet it is the second most forecastable component. While it is true that white noise is slower than an auto-regressive process of order 1 ($AR(1)$) with negative autocorrelation, the latter is still more forecastable. Since we want to reveal intertemporal structure, white noise must be ranked lowest; and ForeCA indeed does so (Fig. 2d).

ForeC 5 and 8 detect the 20 day lag (one trading month), but correlations are too low to achieve much higher forecastability than – simpler and faster – SFA.

In the next example I study quarterly income data, where ForeCA can leverage its nonparametric power and detect important dependencies at various frequencies automatically from the data.

5.2. Classification of US State Economies

I consider quarterly per-capita income growth rates of the “lower 48” from 1982/1 to 2011/4 (last 30 years)

$$g_{j,t} = r_{j,t} - r_{US,t}, \quad j \in \{\text{AL}, \dots, \text{WY}\},$$

where $r_{j,t}$ is the annual growth rate of region j .⁶ Interested in finding similar state economies within the US, we subtract the US baseline. Clustering states with similar economic dynamics can help to decide where to

⁶Publicly available at www.bea.gov/itable.

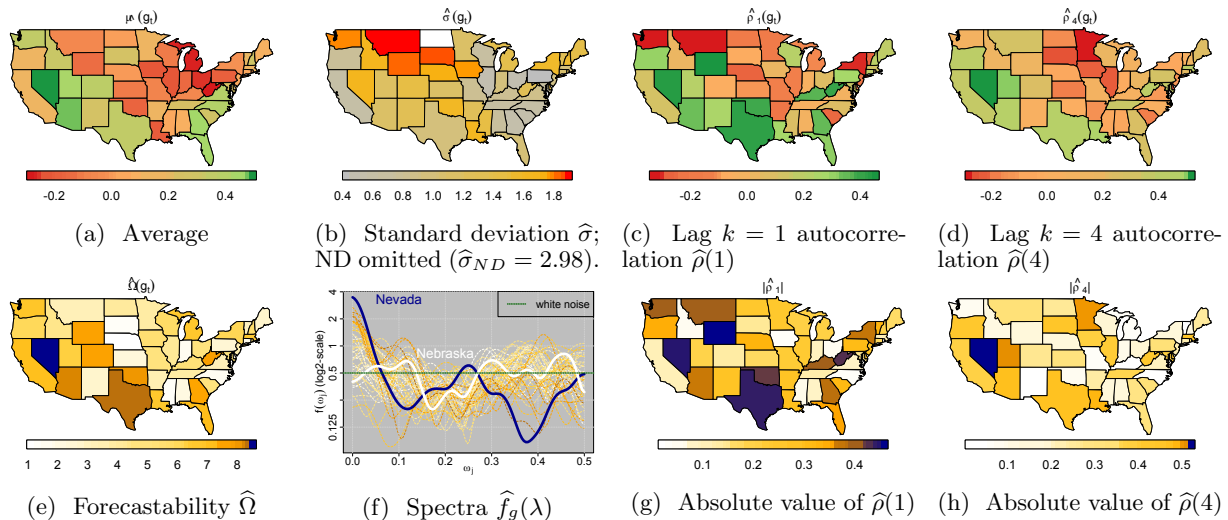


Figure 3. Summary statistics of quarterly income growth rates (in %) from 1982/1 – 2011/4 with respect to US baseline $\hat{\mu}(r_{US,t}) = 1.32\%$, $\hat{\sigma}(r_{US,t}) = 0.92\%$ per quarter; $\hat{\Omega}(r_{US,t}) = 4.86\%$, $\hat{\rho}_1(r_{US,t}) = 0.42$, $\hat{\rho}_4(r_{US,t}) = 0.13$.

provide support when facing difficult economic times. For example, if certain states do not show any important dynamics on a 7-8 year scale – also known as the “business cycle” (Hughes Hallett and Richter, 2008) – then it might be better to support states that are affected by these global economy swings.

The first row of Fig. 3 displays basic summary statistics: sample average, standard deviation, and first and fourth order autocorrelation. The second row give statistics related to forecastability: Fig. 3e shows $\hat{\Omega}$ based on the spectra in Fig. 3f; Fig. 3g shows the absolute lag 1 correlation (analogously for lag 4 in Fig. 3h), since two $AR(1)$ s with a $\pm\phi$ lag 1 coefficient are equivalent in terms of forecasting (compare to SFA ranking in the portfolio example).

The spectral densities of Nevada and Nebraska illustrate the intuitive derivation of $\Omega(x_t)$ from Eq. (10): for Nebraska all frequencies are equally important and it is thus difficult to forecast any better than the sample mean; contrary, Nevada’s income growth rates are mainly driven by a yearly cycle ($\omega_j \approx 0.25$) and low frequencies, thus Nevada is much easier to forecast.

A similar dataset (but annually and for different years) has been analyzed in Dhiral, Kalpakis, Gada, and Puttagunta (2001), who fit $AR(1)$ models to the non-adjusted growth rates $r_{j,t}$ for 25 pre-selected states, and then cluster them in the model space. Although they obtain interpretable results, it is unlikely that US state economies only differ in their lag 1 coefficient. In particular, simple $AR(1)$ models cannot capture the business cycle, which is clearly visible in Fig. 3f (even for the adjusted rates).

Similarly, as SFA maximizes lag 1 correlation, it misses the quarterly cycle. ForeCA does not face this model selection bias, but can find forecastability across all frequencies. In particular, only ForeC 4 detects interesting high frequency signals (Fig. 4b). The most forecastable PCs are PC 5, 4, and 1; interestingly PC 3 is least important for forecasting among all 48 PCs. Also note that ForeCs are more interpretable than SFs or PCs (Figs. 4b - 4d). Particularly, ForeC 1 shows a clear ≈ 25 year period (generation cycle), whereas PC 1 looks somewhat arbitrary. Yet, the associated loadings in Fig. 4a are quite similar.

6. Related Work

Using predictability to separate signals is not new.

In the classic time series literature Box and Tiao (1977) introduced canonical analysis and measure predictive power by the residual variance of fitting vector auto-regression (VAR) models. Recently Matteson and Tsay (2011) propose another DR technique that blends PCA and ICA by separating signals to the extent of fourth moments (but not higher).

Stone (2001) use predictability as a contrast function for blind source separation (BSS). While their approach is similar to ours, it relies on subjective measures of “short” and “long” term moving averages, which are then used to produce actual forecasts.

Much work in BSS (Gomez-Herrero, Rutanen, and Egiazarian, 2010; Li and Adali, 2010), especially ICA, focuses on minimizing entropy rate. The entropy rate $\mathcal{H}(y_t) = \lim_{t \rightarrow \infty} H(y_t | y_{t-1}, y_{t-2}, \dots)$ of a Gaus-

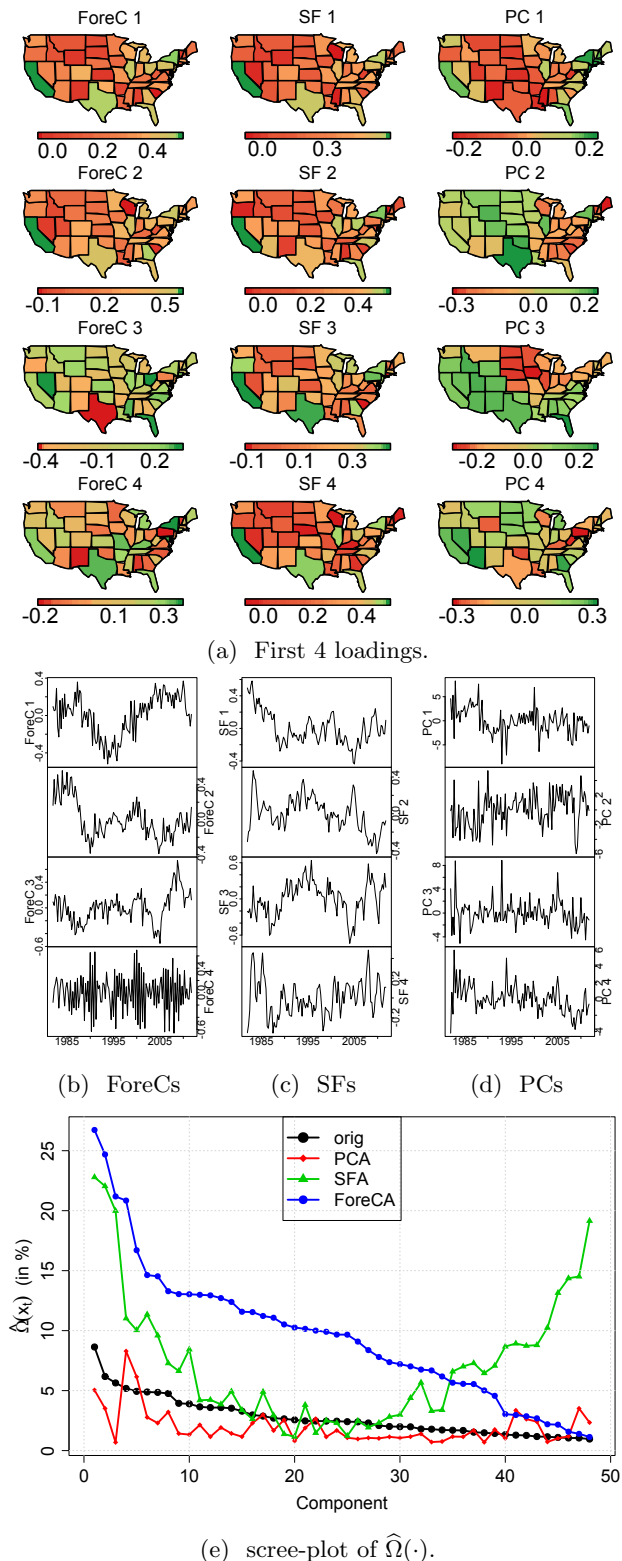


Figure 4. PCA, SFA, and ForeCA on US income data.

sian process is related to the spectrum via (Cover and

Thomas, 1991, p. 417)

$$\mathcal{H}(y_t) = \frac{1}{2} \log 2\pi e + \frac{1}{4\pi} \int_{-\pi}^{\pi} \log S_y(\lambda) d\lambda. \quad (28)$$

However, these approaches require VAR model fits and/or numerical optimization.

On the contrary, the ForeCA measure $\Omega(y_t)$ is based on information-theoretic uncertainty and is an inherent property of the stochastic process y_t . We believe that this makes $\Omega(y_t)$ a more principled measure of forecastability than model-dependent measures. Furthermore, it can be estimated quickly using data-driven, nonparametric techniques.

It is important to point out that *spectral entropy*, i.e., differential entropy of (11), is neither equal nor proportional to the entropy rate in (28). For particular processes they coincide (e.g., for an $AR(1)$; Gibson (1994)), but in general they don't. They measure *different* properties of the signal. Thus ICA algorithms based on entropy rate minimization do not yield the same results as ForeCA. In fact, the ForeCA measure can be used to rank ICs by decreasing forecastability.

Cardoso (2004) gives an excellent account of the intertwined relations between Gaussianity, autocorrelation, and dependence in multivariate time series and their effect on objective functions for BSS. Exactly because of this tangle, we only consider frequency properties of the signal and not entropy rate – since for forecasting the distribution itself is of minor importance compared to the temporal dependence.

7. Discussion

I introduce Forecastable Component Analysis (ForeCA), a new dimension reduction technique for multivariate time series. Contrary to other popular methods – such as PCA or ICA – ForeCA takes temporal dependence into account and actively searches for the most forecastable subspace. ForeCA minimizes the entropy of the spectral density: lower entropy implies a more forecastable signal. The optimization problem has an iterative, yet fast analytic solution, and provably leads to a (local) optimum.

While SFA is a good approximation (maximizing lag 1 correlation), real world signals often have more complex correlation structure. The here proposed ForeCA can automatically detect arbitrary autocorrelation structure using nonparametric estimators. Applications to financial and macro-economic data demonstrate that ForeCA is better than PCA and SFA at finding the most predictable signals, and can also be used for time series classifications.

References

- Box, G. E. P. and G. C. Tiao (1977). A canonical analysis of multiple time series. *Biometrika* 64(2), 355–365.
- Brockwell, P. J. and R. A. Davis (1991). *Time Series: Theory and Methods* (2 ed.). New York, NY: Springer Series in Statistics.
- Cardoso, J.-F. (2004). Dependence, correlation and gaussianity in independent component analysis. *J. Mach. Learn. Res.* 4(7-8), 1177–1203.
- Cover, T. M. and J. Thomas (1991). *Elements of Information Theory*. Wiley.
- Dempster, A. P., N. M. Laird, and D. B. Rubin (1977). Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society B Methodological* 39(1), 1–38.
- Dhiral, K. K., K. Kalpakis, D. Gada, and V. Puttagunta (2001). Distance Measures for Effective Clustering of ARIMA Time-Series. In *Proceedings of the 2001 IEEE International Conference on Data Mining*, pp. 273–280.
- Fessler, J. A. and B. P. Sutton (2003). Nonuniform fast fourier transforms using min-max interpolation. *IEEE Trans. Signal Process* 51, 560–574.
- Fryzlewicz, P., G. P. Nason, and R. von Sachs (2008). A wavelet-Fisz approach to spectrum estimation. *Journal of Time Series Analysis* 29(5), 868–880.
- Gibson, J. (1994). What is the interpretation of spectral entropy? In *Proceedings of IEEE International Symposium on Information Theory, 1994*, pp. 440.
- Gibson, J., S. Stanners, and S. McClellan (1993). Spectral entropy and coefficient rate for speech coding. In *Signals, Systems and Computers, 1993. 1993 Conference Record of The Twenty-Seventh Asilomar Conference on*, pp. 925–929 vol.2.
- Gomez-Herrero, G., K. Ratanen, and K. Egiazarian (2010). Blind source separation by entropy rate minimization. *Signal Processing Letters, IEEE* 17(2), 153–156.
- Hughes Hallett, A. and C. Richter (2008). Have the Eurozone economies converged on a common European cycle? *International Economics and Economic Policy* 5, 71–101.
- Hyvärinen, A. and E. Oja (2000). Independent Component Analysis: Algorithms and Applications. *Neural Networks* 13, 411–430.
- Jacques, L. and P. Vandergheynst (2010). *Compressed Sensing: “When sparsity meets sampling”*, Chapter 23, pp. 507–528. Wiley-Blackwell.
- Jolliffe, I. T. (2002). *Principal Component Analysis* (2 ed.). New York, NY: Springer.
- Lees, J. M. and J. Park (1995). Multiple-Taper Spectral-Analysis - A Stand-Alone C-Subroutine. *Computers & Geosciences* 21(2), 199–236.
- Li, X.-L. and T. Adali (2010). Blind spatiotemporal separation of second and/or higher-order correlated sources by entropy rate minimization. In *Acoustics Speech and Signal Processing (ICASSP), 2010 IEEE International Conference on*, pp. 1934–1937.
- Matteson, D. S. and R. S. Tsay (2011). Dynamic orthogonal components for multivariate time series. *Journal of the American Statistical Association* 106(496), 1450–1463.
- Nuttal, A. H. and G. C. Carter (1982). Spectral Estimation and Lag Using Combined Time Weighting. In *Proceedings of IEEE*, Volume 70, pp. 1111–1125.
- Paninski, L. (2003). Estimation of entropy and mutual information. *Neural Comput.* 15(6), 1191–1253.
- R Development Core Team (2010). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing. ISBN 3-900051-07-0.
- Shannon, C. E. (1948). A Mathematical Theory of Communication. *Bell System Technical Journal* 27, 379–23, 623–656.
- Sricharan, K., R. Raich, and A. Hero (2011). K-nearest neighbor estimation of entropies with confidence. In *Information Theory Proceedings (ISIT), 2011 IEEE International Symposium on*, pp. 1205–1209.
- Stone, J. V. (2001). Blind source separation using temporal predictability. *Neural Comput.* 13(7), 1559–1574.
- Stowell, D. and M. D. Plumbley (2009). Fast Multidimensional Entropy Estimation by k-d Partitioning. *IEEE Signal Processing Letters* 16, 537–540.
- Trobs, M. and G. Heinzel (2006). Improved spectrum estimation from digitized time series on a logarithmic frequency axis. *Measurement* 39(2), 120–129.
- Wiskott, L. and T. J. Sejnowski (2002). Slow Feature Analysis: Unsupervised Learning of Invariances. *Neural computation* 14(4), 715–770.