Contents lists available at ScienceDirect





Pattern Recognition Letters

journal homepage: www.elsevier.com/locate/patrec



Philippe Donnat^a, Gautier Marti^{a,b,*}, Philippe Very^a

^a Hellebore Capital Management, 63 Avenue des Champs-Elysées, Paris 75008, France ^b École Polytechnique, 1 rue Honoré d'Estienne d'Orves, Palaiseau 91120, France

ARTICLE INFO

Article history: Received 27 November 2014 Available online 27 November 2015

Keywords: Financial time series i.i.d. random processes Normalizations and preprocessings Clustering random variables

ABSTRACT

This paper presents a pre-processing and a distance which improve the performance of machine learning algorithms working on independent and identically distributed stochastic processes. We introduce a novel nonparametric approach to represent random variables which splits apart dependency and distribution without losing any information. We also propound an associated metric leveraging this representation and its statistical estimate. Besides experiments on synthetic datasets, the benefits of our contribution is illustrated through the example of clustering financial time series, for instance prices from the credit default swaps market. Results are available on the website http://www.datagrapple.com and an IPython Notebook tutorial is available at http://www.datagrapple.com/Tech for reproducible research.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Machine learning on time series is a booming field and as such plenty of representations, transformations, normalizations, metrics and other divergences are thrown at disposal to the practitioner. A further consequence of the recent advances in time series mining is that it is difficult to have a sober look at the state of the art since many papers state contradictory claims as described in [10]. To be fair, we should mention that when data, pre-processing steps, distances and algorithms are combined together, they have an intricate behavior making it difficult to draw unanimous conclusions especially in a fast-paced environment. Restricting the scope of time series to independent and identically distributed (i.i.d.) stochastic processes, we propound a method which, on the contrary to many of its counterparts, is mathematically grounded with respect to the clustering task defined in Section 1.1. The representation we present in Section 2 exploits a property similar to the seminal result of copula theory, namely Sklar's theorem [36]. This approach leverages the specificities of random variables and this way solves several shortcomings of more classical data pre-processing and distances that will be detailed in Section 1.2. Section 3 is dedicated to experiments on synthetic and real datasets to illustrate the benefits of our method which relies on the hypothesis of i.i.d. sampling of the random variables. Synthetic time series are generated by a simple model yielding correlated random variables following different distributions. The presented approach is also applied to financial time series from the credit default swaps market whose prices dynamics are usually modeled by random walks according to the efficient-market hypothesis [12]. This dataset seems more interesting than stocks as credit default swaps are often considered as a gauge of investors' fear, thus time series are subject to more violent moves and may provide more distributional information than the ones from the stock market. We have made our detailed experiments (cf. Machine Tree on the website www. datagrapple.com) and Python code available (www.datagrapple.com/ Tech) for reproducible research. Finally, we conclude the paper with a discussion on the method and we propound future research directions.

1.1. Motivation and goal of study

Machine learning methodology usually consists in several preprocessing steps aiming at cleaning data and preparing them for being fed to a battery of algorithms. Data scientists have the daunting mission to choose the best possible combination of pre-processing, dissimilarity measure and algorithm to solve the task at hand among a profuse literature. In this article, we provide both a pre-processing and a distance for studying i.i.d. random processes which are compatible with basic machine learning algorithms.

Many statistical distances exist to measure the dissimilarity of two random variables, and therefore two i.i.d. random processes. Such distances can be roughly classified in two families:

^{*} This paper has been recommended for acceptance by Andrea Torsello.

^{*} Corresponding author at: Hellebore Capital Management, 63 Avenue des Champs-Elysées, Paris 75008, France. Tel.: +33155272770; fax: +33155272719. *E-mail address:* gautier.marti@helleborecapital.com (G. Marti).

- distributional distances, for instance [19,33] and [15], which focus on dissimilarity between probability distributions and quantify divergences in marginal behaviors,
- dependence distances, such as the distance correlation or copulabased kernel dependency measures [32], which focus on the joint behaviors of random variables, generally ignoring their distribution properties.

However, we may want to be able to discriminate random variables both on distribution and dependence. This can be motivated, for instance, from the study of financial assets returns: are two perfectly correlated random variables (assets returns), but one being normally distributed and the other one following a heavy-tailed distribution, similar? From risk perspective, the answer is no [18], hence the propounded distance of this paper. We illustrate its benefits through clustering, a machine learning task which primarily relies on the metric space considered (data representation and associated distance). Besides clustering has found application in finance, e.g. [38], which gives us a framework for benchmarking on real data.

Our objective is therefore to obtain a good clustering of random variables based on an appropriate and simple enough distance for being used with basic clustering algorithms, e.g. Ward hierarchical clustering [41], *k*-means++ [1], affinity propagation [14].

By clustering we mean the task of grouping sets of objects in such a way that objects in the same cluster are more similar to each other than those in different clusters. More specifically, a cluster of random variables should gather random variables with common dependence between them and with a common distribution. Two clusters should differ either in the dependency between their random variables or in their distributions.

A good clustering is a partition of the data that must be stable to small perturbations of the dataset. "Stability of some kind is clearly a desirable property of clustering methods" [5]. In the case of random variables, these small perturbations can be obtained from resampling [20,22,27] in the spirit of the bootstrap method since it preserves the statistical properties of the initial sample [11].

Yet, practitioners and researchers pinpoint that state-of-the-art results of clustering methodology applied to financial times series are very sensitive to perturbations [21]. The observed unstability may result from a poor representation of these time series, and thus clusters may not capture all the underlying information.

1.2. Shortcomings of a standard machine learning approach

A naive but often used distance between random variables to measure similarity and to perform clustering is the L_2 distance $\mathbb{E}[(X - Y)^2]$. Yet, this distance is not suited to our task.

Example 1 (Distance L_2 between two Gaussians). Let (X, Y) be a bivariate Gaussian vector, with $X \sim \mathcal{N}(\mu_X, \sigma_X^2)$, $Y \sim \mathcal{N}(\mu_Y, \sigma_Y^2)$ and whose correlation is $\rho(X, Y) \in [-1, 1]$. We obtain $\mathbb{E}[(X - Y)^2] = (\mu_X - \mu_Y)^2 + (\sigma_X - \sigma_Y)^2 + 2\sigma_X\sigma_Y(1 - \rho(X, Y))$. Now, consider the following values for correlation:

- $\rho(X, Y) = 0$, so $\mathbb{E}[(X Y)^2] = (\mu_X \mu_Y)^2 + \sigma_X^2 + \sigma_Y^2$. The two variables are independent (since uncorrelated and jointly normally distributed), thus we must discriminate on distribution information. Assume $\mu_X = \mu_Y$ and $\sigma_X = \sigma_Y$. For $\sigma_X = \sigma_Y \gg 1$, we obtain $\mathbb{E}[(X Y)^2] \gg 1$ instead of the distance 0, expected from comparing two equal Gaussians.
- $\rho(X, Y) = 1$, so $\mathbb{E}[(X Y)^2] = (\mu_X \mu_Y)^2 + (\sigma_X \sigma_Y)^2$. Since the variables are perfectly correlated, we must discriminate on distributions. We actually compare them with a L_2 metric on the mean × standard deviation half-plane. However, this is not an appropriate geometry for comparing two Gaussians [6]. For instance, if $\sigma_X = \sigma_Y = \sigma$, we find $\mathbb{E}[(X Y)^2] = (\mu_X \mu_Y)^2$ for any values of σ . As σ grows, probability attached by the two Gaussians to a



Fig. 1. Probability density functions of Gaussians $\mathcal{N}(-5, 1)$ and $\mathcal{N}(5, 1)$ (in green), Gaussians $\mathcal{N}(-5, 3)$ and $\mathcal{N}(5, 3)$ (in red), and Gaussians $\mathcal{N}(-5, 10)$ and $\mathcal{N}(5, 10)$ (in blue). Green, red and blue Gaussians are equidistant using L_2 geometry on the parameter space (μ, σ). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article).

given interval grows similar (cf. Fig. 1), yet this increasing similarity is not taken into account by this L_2 distance.

 $\mathbb{E}[(X - Y)^2]$ considers both dependence and distribution information of the random variables, but not in a relevant way with respect to our task. Yet, we will benchmark against this distance since other more sophisticated distances on time series such as dynamic time warping [3] and representations such as wavelets [31] or SAX [23] were explicitly designed to handle temporal patterns which are inexistant in i.i.d. random processes.

2. A generic representation for random variables

Our purpose is to introduce a new data representation and a suitable distance which takes into account both distributional proximities and joint behaviors.

2.1. A representation preserving total information

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Ω is the sample space, \mathcal{F} is the σ -algebra of events, and \mathbb{P} is the probability measure. Let \mathcal{V} be the space of all continuous real-valued random variables defined on $(\Omega, \mathcal{F}, \mathbb{P})$. Let \mathcal{U} be the space of random variables following a uniform distribution on [0, 1] and \mathcal{G} be the space of absolutely continuous cumulative distribution functions (cdf).

Definition 1 (The copula transform). Let $X = (X_1, ..., X_N) \in \mathcal{V}^N$ be a random vector with cdfs $G_X = (G_{X_1}, ..., G_{X_N}) \in \mathcal{G}^N$. The random vector $G_X(X) = (G_{X_1}(X_1), ..., G_{X_N}(X_N)) \in \mathcal{U}^N$ is known as the copula transform.

Property 1 (Uniform margins of the copula transform). $G_{X_i}(X_i), 1 \le i \le N$, are uniformly distributed on [0, 1].

Proof.
$$x = G_{X_i}(G_{X_i}^{-1}(x)) = \mathbb{P}(X_i \le G_{X_i}^{-1}(x)) = \mathbb{P}(G_{X_i}(X_i) \le x).$$

We define the following representation of random vectors that actually splits the joint behaviors of the marginal variables from their distributional information.

Definition 2 (dependence \oplus distribution space projection). Let \mathcal{T} be a mapping which transforms $X = (X_1, \ldots, X_N)$ into its generic representation, an element of $\mathcal{U}^N \times \mathcal{G}^N$ representing X, defined as follow

$$\mathcal{T}: \mathcal{V}^N \to \mathcal{U}^N \times \mathcal{G}^N$$

$$X \mapsto (G_X(X), G_X).$$
(1)



Fig. 2. ArcelorMittal and Société générale prices (*T* observations ($X_1^{\epsilon}, X_2^{\epsilon})_{t=1}^{T}$ from (X_1, X_2) $\in \mathcal{V}^2$) are projected on dependence \oplus distribution space; ($G_{X_1}(X_1), G_{X_2}(X_2)$) $\in \mathcal{U}^2$ encode the dependence between X_1 and X_2 (a perfect correlation would be represented by a sharp diagonal on the scatterplot); (G_{X_1}, G_{X_2}) are the margins (their log-densities are displayed above), notice their heavy-tailed exponential distribution (especially for ArcelorMittal).

Property 2. T is a bijection.

Proof. \mathcal{T} is surjective as any element $(U, G) \in \mathcal{U}^N \times \mathcal{G}^N$ has the fiber $G^{-1}(U)$. \mathcal{T} is injective as $(U_1, G_1) = (U_2, G_2)$ *a.s.* in $\mathcal{U}^N \times \mathcal{G}^N$ implies that they have the same cdf $G = G_1 = G_2$ and since $U_1 = U_2$ *a.s.*, it follows that $G^{-1}(U_1) = G^{-1}(U_2)$ *a.s.* \Box

This result replicates the seminal result of copula theory, namely Sklar's theorem [36], which asserts one can split the dependency and distribution apart without losing any information. Fig. 2 illustrates this projection for N = 2.

2.2. A distance between random variables

We leverage the propounded representation to build a suitable yet simple distance between random variables which is invariant under diffeomorphism.

Definition 3 (Distance d_{θ} between two random variables). Let $\theta \in [0, 1]$. Let $(X, Y) \in \mathcal{V}^2$. Let $G = (G_X, G_Y)$, where G_X and G_Y are respectively X and Y marginal cdfs. We define the following distance

$$d_{\theta}^{2}(X,Y) = \theta d_{1}^{2}(G_{X}(X),G_{Y}(Y)) + (1-\theta)d_{0}^{2}(G_{X},G_{Y}),$$
(2)

where

$$d_1^2(G_X(X), G_Y(Y)) = 3\mathbb{E}[|G_X(X) - G_Y(Y)|^2],$$
(3)

and

$$d_0^2(G_X, G_Y) = \frac{1}{2} \int_{\mathbf{R}} \left(\sqrt{\frac{dG_X}{d\lambda}} - \sqrt{\frac{dG_Y}{d\lambda}} \right)^2 d\lambda.$$
(4)

In particular, $d_0 = \sqrt{1 - BC}$ is the Hellinger distance related to the Bhattacharyya (1/2-Chernoff) coefficient *BC* upper bounding the Bayes' classification error. To quantify distribution dissimilarity, d_0 is used rather than the more general α -Chernoff divergences since it satisfies the properties [3,4,5] (significant for practitioners). In addition, d_{θ} can thus be efficiently implemented as a scalar product. $d_1 = \sqrt{(1 - \rho_S)/2}$ is a distance correlation measuring statistical dependence between two random variables, where ρ_S is the Spearman's correlation between *X* and *Y*. Notice that d_1 can be expressed by using the copula *C*: [0, 1]² \rightarrow [0, 1] implicitly defined by the relation $G(X, Y) = C(G_X(X), G_Y(Y))$ since $\rho_S(X, Y) = 12 \int_0^1 \int_0^1 C(u, v) du dv -$ 3 [13].

Example 2 (Distance d_{θ} between two Gaussians). Let (*X*, *Y*) be a bivariate Gaussian vector, with $X \sim \mathcal{N}(\mu_X, \sigma_X^2)$, $Y \sim \mathcal{N}(\mu_Y, \sigma_Y^2)$ and $\rho(X, Y) = \rho$. We obtain,

$$d_{\theta}^{2}(X,Y) = \theta \frac{1-\rho_{S}}{2} + (1-\theta) \left(1 - \sqrt{\frac{2\sigma_{X}\sigma_{Y}}{\sigma_{X}^{2} + \sigma_{Y}^{2}}} e^{-\frac{1}{4}\frac{(\mu_{X}-\mu_{Y})^{2}}{\sigma_{X}^{2} + \sigma_{Y}^{2}}}\right).$$

Remember that for perfectly correlated Gaussians ($\rho = \rho_S = 1$), we want to discriminate on their distributions. We can observe that

- for σ_X, σ_Y → +∞, then d₀(X, Y) → 0, it alleviates a main shortcoming of the basic L₂ distance which is diverging to +∞ in this case;
- if $\mu_X \neq \mu_Y$, for σ_X , $\sigma_Y \rightarrow 0$, then $d_0(X, Y) \rightarrow 1$, its maximum value, i.e. it means that two Gaussians cannot be more remote from each other than two different Dirac delta functions.

We will refer to the use of this distance as the generic parametric representation (GPR) approach. GPR distance is a fast and good proxy for distance d_{θ} when the first two moments μ and σ predominate. Nonetheless, for datasets which contain heavy-tailed distributions, GPR fails to capture this information.

Property 3. Let $\theta \in [0, 1]$. The distance d_{θ} verifies $0 \le d_{\theta} \le 1$.

Proof. Let $\theta \in [0, 1]$. We have

(i) $0 \le d_0 \le 1$, property of the Hellinger distance; (ii) $0 \le d_1 \le 1$, since $-1 \le \rho_S \le 1$.

Finally, by convex combination, $0 \le d_{\theta} \le 1$. \Box

Property 4. For $0 < \theta < 1$, d_{θ} is a metric.

Proof. Let $(X, Y) \in \mathcal{V}^2$. For $0 < \theta < 1$, d_θ is a metric, and the only non-trivial property to verify is the separation axiom

- (i) $X = Y a.s. \Rightarrow d_{\theta}(X, Y) = 0$ $X = Y a.s. \Rightarrow d_1(G_X(X), G_Y(Y)) = d_0(G_X, G_Y) = 0$, and thus $d_{\theta}(X, Y) = 0$,
- (ii) $d_{\theta}(X, Y) = 0 \Rightarrow X = Y a.s.$ $d_{\theta}(X, Y) = 0 \Rightarrow d_1(G_X(X), G_Y(Y)) = 0$ and $d_0(G_X, G_Y) = 0 \Rightarrow$ $G_X(X) = G_Y(Y) a.s.$ and $G_X = G_Y$. Since *G* is absolutely continuous, it follows X = Y a.s.

Notice that for $\theta \in \{0, 1\}$, this property does not hold. Let $U \in \mathcal{V}$, $U \sim \mathcal{U}[0, 1]$. $U \neq 1 - U$ but $d_0(U, 1 - U) = 0$. Let $V \in \mathcal{V}$. $V \neq 2V$ but $d_1(V, 2V) = 0$. \Box

Property 5. Diffeomorphism invariance. Let $h : V \to V$ be a diffeomorphism. Let $(X, Y) \in V^2$. Distance d_{θ} is invariant under diffeomorphism, i.e.

$$d_{\theta}(h(X), h(Y)) = d_{\theta}(X, Y).$$
(5)

Proof. From definition, we have

$$d_0^2(h(X), h(Y)) = 1 - \int_{\mathbf{R}} \sqrt{\frac{dG_{h(X)}}{d\lambda}} \frac{dG_{h(Y)}}{d\lambda} d\lambda$$
(6)

and since

$$\frac{dG_{h(X)}}{d\lambda}(\lambda) = \frac{1}{h'(h^{-1}(\lambda))} \frac{dG_X}{d\lambda} (h^{-1}(\lambda)),$$
(7)

we obtain

$$d_0^2(h(X), h(Y)) = 1 - \int_{\mathbf{R}} \frac{1}{h'(h^{-1}(\lambda))} \sqrt{\frac{dG_X}{d\lambda} \frac{dG_Y}{d\lambda}} (h^{-1}(\lambda)) d\lambda$$
$$= d_0^2(X, Y).$$
(8)

In addition, $\forall x \in \mathbf{R}$, we have

$$G_{h(X)}(h(x)) = \mathbb{P}[h(X) \le h(x)]$$

=
$$\begin{cases} \mathbb{P}[X \le x] = G_X(x), \text{ if } h \text{ increasing} \\ 1 - \mathbb{P}[X \le x] = 1 - G_X(x), \text{ otherwise} \end{cases}$$
(9)

which implies that

$$d_{1}^{2}(h(X), h(Y)) = 3\mathbb{E}\Big[|G_{h(X)}(h(X)) - G_{h(Y)}(h(Y))|^{2}\Big]$$

= $3\mathbb{E}\Big[|G_{X}(X) - G_{Y}(Y)|^{2}\Big]$
= $d_{1}^{2}(X, Y).$ (10)

Finally, we obtain Property 5 by definition of d_{θ} . \Box

Thus, d_{θ} is invariant under monotonic transformations, a desirable property as it ensures to be insensitive to scaling (e.g. choice of units) or measurement scheme (e.g. device, mathematical modeling) of the underlying phenomenon.

2.3. A non-parametric statistical estimation of d_{θ}

To apply the propounded distance d_{θ} on sampled data without parametric assumptions, we have to define its statistical estimate \tilde{d}_{θ} working on realizations of the i.i.d. random variables. Distance d_1 working with continuous uniform distributions can be approximated by normalized rank statistics yielding to discrete uniform distributions, in fact coordinates of the multivariate empirical copula [7] which is a non-parametric estimate converging uniformly toward the underlying copula [8]. Distance d_0 working with densities can be approximated by using its discrete form working on histogram density estimates.

Definition 4 (The empirical copula transform). Let $X^T = (X_1^t, \ldots, X_N^t)$, $t = 1, \ldots, T$, be *T* observations from a random vector $X = (X_1, \ldots, X_N)$ with continuous margins $G_X = (G_{X_1}(X_1), \ldots, G_{X_N}(X_N))$. Since one cannot directly obtain the corresponding copula observations $(G_{X_1}(X_1^t), \ldots, G_{X_N}(X_N^t))$ without knowing a priori G_X , one can instead estimate the *N* empirical margins $G_{X_1}^T(x) = \frac{1}{T} \sum_{t=1}^{T} \mathbf{1}(X_t^t \le x)$ to obtain *T* empirical observations $(G_{X_1}^T(X_1^t), \ldots, G_{X_N}^T(X_N^t))$ which are thus related to normalized rank statistics as $G_{X_1}^T(X_1^t) = X_i^{(t)}/T$, where $X_i^{(t)}$ denotes the rank of observation X_i^t .

Definition 5 (Empirical distance). Let $(X^t)_{t=1}^T$ and $(Y^t)_{t=1}^T$ be *T* realizations of real-valued random variables $X, Y \in \mathcal{V}$ respectively. An empirical distance between realizations of random variables can be defined by

$$\tilde{d}_{\theta}^{2} \left((X^{t})_{t=1}^{T}, (Y^{t})_{t=1}^{T} \right) \stackrel{a.s.}{=} \theta \tilde{d}_{1}^{2} + (1-\theta) \tilde{d}_{0}^{2},$$
(11)

where

$$\tilde{d}_1^2 = \frac{3}{T(T^2 - 1)} \sum_{t=1}^{I} \left(X^{(t)} - Y^{(t)} \right)^2$$
(12)

and

$$\tilde{d}_{0}^{2} = \frac{1}{2} \sum_{k=-\infty}^{+\infty} \left(\sqrt{g_{X}^{h}(hk)} - \sqrt{g_{Y}^{h}(hk)} \right)^{2},$$
(13)

h being here a suitable bandwidth, and $g_X^h(x) = \frac{1}{T} \sum_{t=1}^T \mathbf{1}(\lfloor \frac{x}{h} \rfloor h \le X^t < (\lfloor \frac{x}{h} \rfloor + 1)h)$ being a density histogram estimating pdf g_X from $(X^t)_{t=1}^T$, *T* realizations of random variable $X \in \mathcal{V}$.

We will refer henceforth to this distance and its use as the generic non-parametric representation (GNPR) approach. To use effectively d_{θ} and its statistical estimate, it boils down to select a particular value for θ . We suggest here an exploratory approach where one can test (i) distribution information ($\theta = 0$), (ii) dependence information $(\theta = 1)$, and (iii) a mix of both information ($\theta = 0.5$). Ideally, θ should reflect the balance of dependence and distribution information in the data. In a supervised setting, one could select an estimate θ of the right balance θ^* optimizing some loss function by techniques such as cross-validation. Yet, the lack of a clear loss function makes the estimation of θ^* difficult in an unsupervised setting. For clustering, many authors [20,26,34,35] suggest stability as a tool for parameter selection. But, [2] warn against its irrelevant use for this purpose. Besides, we already use stability for clustering validation and we want to avoid overfitting. Finally, we think that finding an optimal tradeoff θ^* is important for accelerating the rate of convergence toward the underlying ground truth when working with finite and possibly small samples, but ultimately lose its importance asymptotically as soon as $0 < \theta < 1$.

3. Experiments and applications

3.1. Synthetic datasets

We propose the following model for testing the efficiency of the GNPR approach: *N* time series of length *T* which are subdivided into *K* correlation clusters themselves subdivided into *D* distribution clusters.

Let $(Y_k)_{k=1}^K$, be *K* i.i.d. random variables. Let $p, D \in \mathbb{N}$. Let N = pKD. Let $(Z_d^i)_{d=1}^D$, $1 \le i \le N$, be independent random variables. For $1 \le i \le N$, we define

$$X_{i} = \sum_{k=1}^{K} \beta_{k,i} Y_{k} + \sum_{d=1}^{D} \alpha_{d,i} Z_{d}^{i},$$
(14)

where

(a)
$$\alpha_{d,i} = 1$$
, if $i \equiv d - 1 \pmod{D}$, 0 otherwise;
(b) $\beta \in [0, 1]$,
(c) $\beta_{k,i} = \beta$, if $\lceil iK/N \rceil = k$, 0 otherwise.

 $(X_i)_{i=1}^N$ are partitioned into Q = KD clusters of p random variables each. Playing with the model parameters, we define in Table 1 some interesting test case datasets to study distribution clustering, dependence clustering or a mix of both. We use the following notations as a shorthand

(a)
$$\mathcal{L} := \text{Laplace}(0, 1/\sqrt{2})$$

(b) $\mathcal{S} := \text{t-distribution}(3)/\sqrt{3}$

Since \mathcal{L} and \mathcal{S} have both a mean of 0 and a variance of 1, GPR cannot find any difference between them, but GNPR can discriminate on higher moments as it can be seen in Fig. 3.

3.2. Performance of clustering using GNPR

We empirically show that the GNPR approach achieves better results than others using common distances regardless of the algorithm used on the defined test cases A, B and C described in Table 1. Test case A illustrates datasets containing only distribution information:

Table 1
Model parameters for some interesting test case datasets.

Clustering	Dataset	Ν	Т	Q	Κ	β	Y_k	Z_1^i	Z_2^i	Z_3^i	Z_4^i
Distribution	А	200	5000	4	1	0	$\mathcal{N}(0,1)$	$\mathcal{N}(0,1)$	L	S	$\mathcal{N}(0,2)$
Dependence	В	200	5000	10	10	0.1	S	S	S	S	S
Mix	С	200	5000	10	5	0.1	$\mathcal{N}(0,1)$	$\mathcal{N}(0,1)$	S	$\mathcal{N}(0,1)$	S
	G	32, , 640	$10,\ldots,2000$	32	8	0.1	$\mathcal{N}(0,1)$	$\mathcal{N}(0,1)$	$\mathcal{N}(0,2)$	\mathcal{L}	S



Fig. 3. GPR and GNPR distance matrices. Both GPR and GNPR highlight the 5 correlation clusters ($\theta = 1$), but only GNPR finds the 2 distributions (S and L) subdividing them ($\theta = 0$). Finally, by combining both information GNPR ($\theta = 0.5$) can highlight the 10 original clusters, while GPR ($\theta = 0.5$) simply adds noise on the correlation distance matrix it recovers.



Fig. 4. Distance matrices obtained on dataset C using distance correlation, L_2 distance, GPR and GNPR. None but GNPR highlights the 10 original clusters which appear on its diagonal.

there are 4 clusters of distributions. Test case B illustrates datasets containing only dependence information: there are 10 clusters of correlation between random variables which are heavy-tailed. Test case C illustrates datasets containing both information: it consists in 10 clusters composed of 5 correlation clusters and each of them is divided into 2 distribution clusters. Using scikit-learn implementation [30], we apply 3 clustering algorithms with different paradigms: a

Table 2 Comparison of distance correlation, L_2 distance, GPR and GNPR: GNPR approach improves clustering performance.

Algo.	Distance	Adjusted rand index						
		A	В	С				
HC-AL	$(1 - \rho)/2$	0.00 ± 0.01	0.99 ± 0.01	0.56 ± 0.01				
	$\mathbb{E}[(X - Y)^2]$	0.00 ± 0.00	0.09 ± 0.12	0.55 ± 0.05				
	$\operatorname{GPR} \theta = 0$	$\textbf{0.34} \pm \textbf{0.01}$	0.01 ± 0.01	0.06 ± 0.02				
	$\operatorname{GPR} \theta = 1$	0.00 ± 0.01	0.99 ± 0.01	0.56 ± 0.01				
	GPR $\theta = .5$	$\textbf{0.34} \pm \textbf{0.01}$	0.59 ± 0.12	0.57 ± 0.01				
	$\operatorname{GNPR} \theta = 0$	1	$\textbf{0.00} \pm \textbf{0.00}$	0.17 ± 0.00				
	$\operatorname{GNPR} \theta = 1$	0.00 ± 0.00	1	0.57 ± 0.00				
	GNPR $\theta = .5$	0.99 ± 0.01	0.25 ± 0.20	0.95 ± 0.08				
KM++	$(1 - \rho)/2$	0.00 ± 0.01	0.60 ± 0.20	0.46 ± 0.05				
	$\mathbb{E}[(X - Y)^2]$	0.00 ± 0.00	0.34 ± 0.11	0.48 ± 0.09				
	$\operatorname{GPR} \theta = 0$	0.41 ± 0.03	0.01 ± 0.01	0.06 ± 0.02				
	$\operatorname{GPR} \theta = 1$	0.00 ± 0.00	0.45 ± 0.11	0.43 ± 0.09				
	GPR $\theta = .5$	0.27 ± 0.05	0.51 ± 0.14	0.48 ± 0.06				
	$\operatorname{GNPR} \theta = 0$	0.96 ± 0.11	$\textbf{0.00} \pm \textbf{0.01}$	0.14 ± 0.02				
	$\operatorname{GNPR} \theta = 1$	0.00 ± 0.01	0.65 ± 0.13	$\textbf{0.53} \pm \textbf{0.02}$				
	$\text{GNPR}\theta = .5$	0.72 ± 0.13	0.21 ± 0.07	0.64 ± 0.10				
AP	$(1 - \rho)/2$	0.00 ± 0.00	0.99 ± 0.07	0.48 ± 0.02				
	$\mathbb{E}[(X - Y)^2]$	0.14 ± 0.03	0.94 ± 0.02	0.59 ± 0.00				
	$\operatorname{GPR} \theta = 0$	0.25 ± 0.08	0.01 ± 0.01	0.05 ± 0.02				
	$\operatorname{GPR} \theta = 1$	0.00 ± 0.01	0.99 ± 0.01	0.48 ± 0.02				
	GPR $\theta = .5$	0.06 ± 0.00	0.80 ± 0.10	0.52 ± 0.02				
	$\operatorname{GNPR} \theta = 0$	1	0.00 ± 0.00	0.18 ± 0.01				
	$\operatorname{GNPR} \theta = 1$	0.00 ± 0.01	1	0.59 ± 0.00				
	$\mathrm{GNPR}\theta=.5$	$\textbf{0.39} \pm \textbf{0.02}$	$\textbf{0.39} \pm \textbf{0.11}$	1				

hierarchical clustering using average linkage (HC-AL), *k*-means++ (KM++), and affinity propagation (AP). Experiment results are reported in Table 2. GNPR performance is due to its proper representation (cf. Fig. 4). Finally, we have noticed increasing precision of clustering using GNPR as time *T* grows to infinity, all other parameters being fixed. The number of time series *N* seems rather uninformative as illustrated in Fig. 5 (left) which plots ARI [16] between computed clustering and ground-truth of dataset G as an heatmap for varying *N*



Fig. 5. Empirical consistency of clustering using GNPR as $T \rightarrow \infty$.

and *T*. Fig. 5 (right) shows the convergence to the true clustering as a function of *T*.

3.3. Application to financial time series clustering

3.3.1. Clustering assets: a (too) strong focus on correlation

It has been noticed that straightfoward approaches automatically discover sector and industries [24]. Since detected patterns are blatantly correlation-flavoured, it prompted econophysicists to focus on correlations, hierarchies and networks [40] from the Minimum Spanning Tree and its associated clustering algorithm the Single Linkage to the state of the art [28] exploiting the topological properties of the Planar Maximally Filtered Graph [39] and its associated algorithm the Directed Bubble Hierarchical Tree (DBHT) technique [37]. In practice, econophysicists consider the assets log returns and compute their correlation matrix. The correlation matrix is then filtered thanks to a clustering of the correlation-network [9] built using similarity and dissimilarity matrices which are derived from the correlation one by convenient ad hoc transformations. Clustering these correlationbased networks [29] aims at filtering the correlation matrix for standard portfolio optimization [38]. Yet, adopting similar approaches only allow to retrieve information given by assets co-movements and nothing about the specificities of their returns behavior, whereas we may also want to distinguish assets by their returns distribution. For example, we are interested to know whether they undergo fat tails, and to which extent.

3.3.2. Clustering credit default swaps

We apply the GNPR approach on financial time series, namely daily credit default swap [17] (CDS) prices. We consider the N = 500 most actively traded CDS according to DTCC (http://www.dtcc.com/). For each CDS, we have T = 2300 observations corresponding to historical daily prices over the last 9 years, amounting for more than one million data points. Since credit default swaps are traded overthe-counter, closing time for fixing prices can be arbitrarily chosen, here 5pm GMT, i.e. after the London Stock Exchange trading session. This synchronous fixing of CDS prices avoids spurious correlations arising from different closing times. For example, the use of close-to-close stock prices artificially overestimates intra-market correlation and underestimates inter-market dependence since they have different trading hours [25]. These CDS time series can be consulted on the web portal http://www.datagrapple.com/.

Assuming that CDS prices $(P^t)_{t \ge 1}$ follow random walks, their increments $\Delta P^t = P^t - P^{t-1}$ are i.i.d. random variables, and therefore the GNPR approach can be applied to the time series of prices variations, i.e. on data $(\Delta P_1^t, \ldots, \Delta P_N^t)$, $t = 1, \ldots, T$. Thus, for aggregating CDS prices time series, we use a clustering algorithm (for instance, Ward's method [41]) based on the GNPR distance matrices between their variations.

Using GNPR $\theta = 0$, we look for distribution information in our CDS dataset. We observe that clustering based on the GNPR $\theta = 0$ distance matrix yields 4 clusters which fit precisely the multi-modal empirical distribution of standard deviations as can be seen in Fig. 6. For GNPR $\theta = 1$, we display in Fig. 7 the rank correlation distance matrix obtained. We can notice its hierarchical structure already described in many papers, e.g. [4,24], focusing on stock markets. There is information in distribution and in correlation, thus taking into account both information, i.e. using GNPR θ = 0.5, should lead to a meaningful clustering. We verify this claim by using stability as a criterion for validation. Practically, we consider even and odd trading days and perform two independent clusterings, one on even days and the other one on odd days. We should obtain the same partitions. In Fig. 8, we display the partitions obtained using the GNPR $\theta = 0.5$ approach next to the ones obtained by applying a L_2 distance on prices returns. We find that GNPR clustering is more stable than L_2 on returns cluster-



Fig. 6. Standard Deviation Histogram. The 4 clusters found using GNPR $\theta = 0$ represented by the 4 colors fit precisely the multi-modal distribution of standard deviations.



Fig. 7. Centered Rank Correlation Distance Matrix. GNPR $\theta = 1$ exhibits a hierarchical structure of correlations: first level consists in Europe, Japan and US; second level corresponds to credit quality (investment grade or high yield); third level to industrial sectors.

ing. Moreover, clusters obtained from GNPR are more homogeneous in size.

To conclude on the experiments, we have highlighted through clustering that the presented approach leveraging dependence and distribution information leads to better results: finer partitions on synthetic test cases and more stable partitions on financial time series.

4. Discussion

In this paper, we have exposed a novel representation of random variables which could lead to improvements in applying machine learning techniques on time series describing underlying i.i.d. stochastic processes. We have empirically shown its relevance to deal with random walks and financial time series. We have led a large scale experiment on the credit derivatives market notorious for not having Gaussian but heavy-tailed returns, first results are available on website www.datagrapple.com. We also intend to lead such clustering experiments for testing applicability of the method to areas outside finance. On the theoretical side, we plan to improve the aggregation of the correlation and distribution part by using elements of information geometry theory and to study the consistency property of our method.



Fig. 8. Better clustering stability using the GNPR approach: GNPR $\theta = 0.5$ achieves ARI = 0.85; L_2 on returns achieves ARI 0.64; The two leftmost partitions built from GNPR on the odd/even trading days sampling look similar: only a few CDS are switching from clusters; The two rightmost partitions built using a L_2 on returns display very inhomogeneous (odd-2,3,9 vs. odd-4,14,15) and unstable (even-1 splitting into odd-3 and odd-2) clusters.

Acknowledgments

We thank Frank Nielsen, the anonymous reviewers, and our colleagues at Hellebore Capital Management for giving feedback and proofreading the article.

References

- D. Arthur, S. Vassilvitskii, k-means++: the advantages of careful seeding, in: SODA '07: Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms, 2007, pp. 1027–1035.
- [2] S. Ben-David, U. Von Luxburg, D. Pál, A sober look at clustering stability, Learn. Theory (2006).
- [3] D. Berndt, J. Clifford, Using dynamic time warping to find patterns in time series, KDD workshop 10 (1994) 359–370.
- [4] G. Brida, A. Risso, Hierarchical structure of the german stock market, Expert Syst. Appl. 37 (2010) 3846–3852.
- [5] G. Carlsson, F. Mémoli, Characterization, stability and convergence of hierarchical clustering methods, J. Mach. Learn. Res. 11 (2010) 1425–1470.
- [6] S. Costa, S. Santos, J. Strapasson, Fisher information distance: a geometrical reading, Discr. Appl. Math. 197 (2014) 59–69.
- [7] P. Deheuvels, La fonction de dépendance empirique et ses propriétés. un test non paramétrique d'indépendance, Acad. R. Belg, Bull. Cl. Sci. 65 (5) (1979) 274–292.
- [8] P. Deheuvels, An asymptotic decomposition for multivariate distribution-free tests of independence, J. Multivar. Anal. 11 (1981) 102–113.
- [9] T. Di Matteo, F. Pozzi, T. Aste, The use of dynamical networks to detect the hierarchical organization of financial market sectors, Eur. Phys. J. B - Condens. Matter Complex Syst. 73 (2010) 3–11.
- [10] H. Ding, G. Trajcevski, P. Scheuermann, X. Wang, E. Keogh, Querying and mining of time series data: experimental comparison of representations and distance measures 1 (2008) 1542–1552.
- [11] B. Efron, Bootstrap methods: another look at the jackknife, Ann. Stat. (1979) 1–26.
- [12] E.F. Fama, The behavior of stock-market prices, J. Bus. 38 (1965) 34–105.
- [13] G. Fredricks, R. Nelsen, On the relationship between spearman's rho and kendall's tau for pairs of continuous random variables, J. Stat. Plann. Inference 137 (2007) 2143–2150.
- [14] B. Frey, D. Dueck, Clustering by passing messages between data points, Science 315 (2007) 972–976.

- [15] K. Henderson, B. Gallagher, T. Eliassi-Rad, EP-MEANS: an efficient nonparametric clustering of empirical probability distributions, in: SAC'15 Proceedings of the 30th Annual ACM Symposium on Applied Computing, 2015, pp. 893–900.
- [16] L. Hubert, P. Arabie, Comparing partitions, J. Classif. 2 (1985) 193-218.
- [17] J.C. Hull, Options, Futures, and Other Derivatives, Pearson Education, 2006.
- [18] B. Kelly, H. Jiang, Tail risk and asset prices, Rev. Finan. Stud. 27 (10) (2014) 2841– 2871
- [19] A. Khaleghi, D. Ryabko, J. Mary, P. Preux, Online clustering of processes 22 (2012) 601–609.
- [20] T. Lange, V. Roth, M. Braun, J. Buhmann, Stability-based model selection, Neural Comput. 16 (2004) 1299–1323.
- [21] V. Lemieux, P. Rahmdel, R. Walker, B.L. Wong, M. Flood, Clustering techniques and their effect on portfolio formation and risk analysis, in: DSMM'14 Proceedings of the International Workshop on Data Science for Macro-Modeling, 2014, pp. 1– 6
- [22] E. Levine, E. Domany, Resampling method for unsupervised estimation of cluster validity, Neural Comput. 13 (2001) 2573–2593.
- [23] J. Lin, E. Keogh, S. Lonardi, B. Chiu, A symbolic representation of time series, with implications for streaming algorithms, Proceedings of the 8th ACM SIGMOD workshop on Research issues in data mining and knowledge discovery (2003) 2– 11
- [24] R.N. Mantegna, Hierarchical structure in financial markets, Eur. Phys. J. B 11 (1999) 193–197.
- [25] M. Martens, S. Poon, Returns synchronization and daily correlation dynamics between international stock markets, J. Bank. Finan. 25 (2001) 1805–1827.
- [26] N. Meinshausen, P. Bühlmann, Stability selection, J. R. Stat. Soc.: Ser. B (Stat. Methodol.) 72 (2010) 417–473.
- [27] S. Monti, P. Tamayo, J. Mesirov, T. Golub, Consensus clustering: a resamplingbased method for class discovery and visualization of gene expression microarray data, Mach. Learn. 52 (2003) 91–118.
- [28] N. Musmeci, T. Aste, T. Di Matteo, Relation between financial market structure and the real economy: comparison between clustering methods, PLoS ONE 10 (3) (2015) e0116201.
- [29] J.-P. Onnela, K. Kaski, J. Kertész, Clustering and information in correlation based financial networks, Eur. Phys. J. B-Condens. Matter Complex Syst. 38 (2004) 353– 362.
- [30] F. Pedregosa, G. Varoquaux, Scikit-learn: machine learning in Python, J. Mach. Learn. Res. 12 (2011) 2825–2830.
- [31] D.B. Percival, A.T. Walden, Wavelet methods for time series analysis (Vol. 4), Cambridge University Press, 2006.

- [32] B. Póczos, Z. Ghahramani, J. Schneider, Copula-based Kernel Dependency Measures, in: Proceedings of the 29th International Conference on Machine Learning (ICML-12), 2012, pp. 775-782.

- (ICML-12), 2012, pp. 175-782.
 [33] D. Ryabko, Clustering processes, ICML (2010).
 [34] O. Shamir, N. Tishby, Cluster Stability for Finite Samples, NIPS (2007).
 [35] O. Shamir, N. Tishby, Model selection and stability in k-means clustering, Learn. Theory (2008).
- [36] A. Sklar, Fonctions de répartition à n dimensions et leurs marges, Publ. Inst. Statist. Univ. Paris (in French) 8 (1959) 229–231.
- [37] W.M. Song, D. Matteo, T. Aste, Hierarchical information clustering by means of topologically embedded graphs, PLoS ONE 7 (2012) e31929+.
- [38] V. Tola, F. Lillo, M. Gallegati, R.N. Mantegna, Cluster analysis for portfolio opti-mization, J. Econ. Dyn. Control 32 (2008) 235–258.
- [39] M. Tumminello, T. Aste, T. Di Matteo, R.N. Mantegna, A tool for filtering informa-[40] M. Tumminello, F. Matter, F. Di Matter, K. Martegia, A teor for intering information in complex systems, Proc. Natl. Acad. Sci. USA 102 (2005) 10421–10426.
 [40] M. Tumminello, F. Lillo, Mantegna, Correlation, hierarchies, and networks in fi-
- nancial markets, J. Econ. Behav. Org. 75 (1) (2010) 40–58.
- [41] J.H. Ward, Hierarchical grouping to optimize an objective function, J. Am. Stat. Assoc. 58 (1963) 236–244.