Measuring the Influence of Observations in HMMs through the Kullback-Leibler Distance

Vittorio Perduca and Gregory Nuel

Abstract—We measure the influence of individual observations on the sequence of the hidden states of the Hidden Markov Model (HMM) by means of the Kullback-Leibler distance (KLD). Namely, we consider the KLD between the conditional distribution of the hidden states' chain given the complete sequence of observations and the conditional distribution of the hidden chain given all the observations but the one under consideration. We introduce a linear complexity algorithm for computing the influence of all the observations. As an illustration, we investigate the application of our algorithm to the problem of detecting meaningful observations in HMM data series.

Index Terms—Hidden Markov models, relative entropy, forward-backward algorithm, outlier detection, local outlier factor

I. INTRODUCTION

T HE Hidden Markov Model (HMM) is a standard tool in many applications, including signal processing and speech recognition [1], [2], [3] and computational biology [4]. In a typical HMM, let $S_{1:n} = (S_1, \ldots, S_n)$ be the Markov sequence of hidden variables (or states) and $X_{1:n} =$ (X_1, \ldots, X_n) the sequence of observation variables¹. In this letter we address the problem of measuring the influence of an observation $X_j = x_j$ on the distribution of the hidden sequence $S_{1:n}$.

We start by fixing notation. For simplicity's sake, we consider *homogeneous* HMMs and denote the parameters of the model with $\mathbb{P}(X_i = x | S_i = s) = \beta(s, x)$ (*emissions*), $\mathbb{P}(S_i = s | S_{i-1} = r) = \alpha(r, s)$ (transitions) and $\mathbb{P}(S_1 = s) = \gamma(s)$. The model is fully specified by the conditional dependencies among the variables depicted in Fig. 1 which determine the following factorization of the joint probability distribution

$$\mathbb{P}(X_{1:n} = x_{1:n}, S_{1:n} = s_{1:n}) = \gamma(s) \prod_{i=2}^{n} \alpha(s_{i-1}, s_i) \prod_{i=1}^{n} \beta(s_i, x_i),$$

where s_i, x_i are taken in the sets of all possible outcomes of S_i and X_i (for continuous variables simply replace probabilities with densities and sums with integrals). For simplicity of notation, in most equations we omit to write explicitly the outcomes of the variables.

An important inference problem in HMMs is computing the conditional (posterior) distribution of the hidden sequence given an *evidence*. In standard applications, the evidence is a complete instantiation of the observable sequence, $\mathcal{E} = \{X_{1:n} = x_{1:n}\}$ for some $x_{1:n}$. For a fixed $j \in \{1, \ldots, n\}$, we denote \mathcal{E}_{-j} the evidence $\{X_{-j} = x_{-j}\}$, where X_{-j} denotes the sequence of all the observation variables except X_j .

Our suggestion for measuring the influence of an observation $X_j = x_j$ is based on the following question: what is the contribution of $X_j = x_j$ to the posterior distribution $S_{1:n}|\{X_{1:n} = x_{1:n}\}$ of the hidden sequence given the complete sequence of observations? That is, how *dissimilar* are the posterior distributions $\mathbb{P}(S_{1:n}|\mathcal{E})$ and $\mathbb{P}(S_{1:n}|\mathcal{E}_{-j})$? The more these two posterior distributions are *distant*, the more $X_j = x_j$ must be influential.

The Kullback-Leibler distance (KLD) (or relative entropy) arises in many applications as an appropriate measurement of the distance between two probability distributions [5], [6]. Following [7] and [8] (in the context of linear regression), we suggest to measure the influence of $X_i = x_i$ through the KLD

$$K_j := \sum_{S_{1:n}} \mathbb{P}(S_{1:n} | \mathcal{E}_{-j}) \log \frac{\mathbb{P}(S_{1:n} | \mathcal{E}_{-j})}{\mathbb{P}(S_{1:n} | \mathcal{E})}.$$

By definition, K_j measures the influence of observation x_j on the posterior distribution of the hidden states rather than on the parameter estimate. K_j is therefore an appropriate influence measure when the quantity of interest is the posterior distribution as it is often the case in practical applications such as speech recognition [3], data segmentation [9], bioinformatics [10], genetics [11].

In this letter we address the problem of computing efficiently the vector $(K_j)_{j=1,...,n}$ of all the KL distances, one for each observation. To the best of our knowledge, the computation of the KLD between the posterior distributions of the hidden sequence of an HMM conditioned on two distinct evidences was not studied before, the main efforts being rather aimed at computing efficiently the KLD between the distributions of the observation sequence of an HMM with respect to two distinct sets of parameters [6], [12], [13].

A straightforward computation of $(K_j)_{j=1,...,n}$ based on the standard forward-backward algorithm for HMMs leads to a quadratic complexity in the number of observations; our main contribution is a linear time algorithm based on simple recursive formulae.

As an illustration, we apply our algorithm to a time series of temperature changes and discuss the practical interest of the suggested influence measure.

V. Perduca and G. Nuel are with the Laboratoire MAP5, Université Paris Descartes and CNRS, Sorbonne Paris Cité, France, email: {vittorio.perduca, gregory.nuel}@parisdescartes.fr.

¹In the absence of a widespread standard notation, symbols denoting HMM variables vary from author to author.



Fig. 1: HMM topology. S_j : hidden variable, X_j : observed variable.

II. COMPUTATION OF THE INFLUENCE MEASURE

We start by recalling that the posterior distribution $\mathbb{P}(S_{1:n}|\mathcal{E})$ of the hidden sequence given the standard evidence is an heterogeneous Markov sequence whose transition probabilities are computed in $\mathcal{O}(nm^2)$ steps with the standard forward-backward algorithm, where *m* is the is the number of possible outcomes of each hidden variable, [2]. The forward and backward quantities are defined as $F_i(s) := \mathbb{P}(X_{1:i}, S_i = s)$, $B_i(s) = \mathbb{P}(X_{i+1:n}|S_i = s)$ and are computed recursively with Eqs. (1) and (2).

Similarly to $\mathbb{P}(S_{1:n}|\mathcal{E})$, the computation of $\mathbb{P}(S_{1:n}|\mathcal{E}_{-j})$, for a fixed j, requires $\mathcal{O}(nm^2)$ steps: it is easy to adapt the forward-backward algorithm to \mathcal{E}_{-j} by simply marginalizing out the variable X_j in all the formulae and propagating the new forward and backward quantities thus obtained.

It is straightforward to compute recursively the KLD between two heterogeneous Markov Chain; as a consequence, a direct approach based on standard recursions leads to a $\mathcal{O}(nm^2)$ time complexity for computing K_j for a fixed j. However, the required marginalization of the forward and backward quantities depends on the fixed j and therefore it is necessary to compute a distinct set of forward and backward quantities for each j. As a consequence, the resulting complexity for computing the vector $(K_j)_{j=1,...,n}$ is $\mathcal{O}(n^2m^2)$.

Our principal contribution are new recursive formulae that reduce this complexity to $\mathcal{O}(nm^2)$. We start with two technical lemmas that lead to our original algorithm.

Lemma 1: For an arbitrary fixed $j \in \{1, \ldots, n\}$:

$$K_j = \sum_{S_j} \mathbb{P}(S_j | \mathcal{E}_{-j}) \log \frac{\mathbb{P}(S_j | \mathcal{E}_{-j})}{\mathbb{P}(S_j | \mathcal{E})}.$$

Proof: We start by observing that the following factorizations hold:

$$\mathbb{P}(S_{1:n}|\mathcal{E}_{-j}) = \mathbb{P}(S_j|\mathcal{E}_{-j})\mathbb{P}(S_{1:j-1}|S_j,\mathcal{E}_{-j})\mathbb{P}(S_{j+1:n}|S_j,\mathcal{E}_{-j})$$

and

$$\mathbb{P}(S_{1:n}|\mathcal{E}) = \mathbb{P}(S_j|\mathcal{E})\mathbb{P}(S_{1:j-1}|S_j,\mathcal{E})\mathbb{P}(S_{j+1:n}|S_j,\mathcal{E}).$$

The key point is that in the last equation we have

$$\mathbb{P}(S_{j+1:n}|S_j,\mathcal{E}) = \mathbb{P}(S_{j+1:n}|S_j,\mathcal{E}_{-j})$$
$$\mathbb{P}(S_{1:j-1}|S_j,\mathcal{E}) = \mathbb{P}(S_{1:j-1}|S_j,\mathcal{E}_{-j})$$

because $S_{j+1:n}$ and X_j are conditionally independent given S_j , and $S_{1:j-1}$ and X_j are conditionally independent given S_j , see Fig. 1. Then $K_j = \sum_{S_j} \mathbb{P}(S_j | \mathcal{E}_{-j}) \log \frac{\mathbb{P}(S_j | \mathcal{E}_{-j})}{\mathbb{P}(S_j | \mathcal{E})} \times \sum_{S_{1:j-1}} \mathbb{P}(S_{1:j-1} | S_j, \mathcal{E}_{-j}) \sum_{S_{j+1:n}} \mathbb{P}(S_{j+1:n} | S_j, \mathcal{E}_{-j})$.

As a consequence of this lemma, the key for computing efficiently $(K_j)_{j=1,...,n}$ is an efficient computation of the factors $\mathbb{P}(S_j|\mathcal{E})$ and $\mathbb{P}(S_j|\mathcal{E}_{-j})$ for all j = 1,...,n. For a given j, $\mathbb{P}(S_j|\mathcal{E})$ can be computed in $\mathcal{O}(nm^2)$ steps using the standard forward-backward algorithm: $\mathbb{P}(S_j = s, \mathcal{E}) = F_j(s)B_j(s)$ and hence $\mathbb{P}(S_j = s|\mathcal{E}) \propto F_j(s)B_j(s)$, where the standard forward and backward quantities are computed recursively with

$$F_i(s) = \sum_r F_{i-1}(r)\alpha(r,s)\beta(s,x_i) \tag{1}$$

and

$$B_{i-1}(r) = \sum_{s} \alpha(r,s)\beta(s,x_i)B_i(s).$$
⁽²⁾

We show a similar result for $\mathbb{P}(S_j|\mathcal{E}_{-j})$:

Lemma 2: For an arbitrary fixed $j \in \{1, \ldots, n\}$:

$$\mathbb{P}(S_j = s, \mathcal{E}_{-j}) = F_j^*(s)B_j(s),$$

where B_j is the standard backward quantity and F_j^* is computed recursively from the standard forward quantities with

$$F_i^*(s) = \sum_r F_{i-1}(r)\alpha(r,s)$$
 for $i = 2, \dots, n,$ (3)

with $F_1^*(s) = \gamma(s)$. Moreover the time complexity for computing $\mathbb{P}(S_j = s | \mathcal{E}_{-j}) \propto F_j^*(s) B_j(s)$ for all j = 1, ..., n is $\mathcal{O}(nm^2)$.

Proof: For a given j, we have $\mathbb{P}(S_j = s, \mathcal{E}_{-j}) =$

$$\sum_{y} \mathbb{P}(S_j = s, \mathcal{E}_{-j}, X_j = y) = \sum_{y} \mathbb{P}(S_j = s, \mathcal{E}_y),$$

where \mathcal{E}_y is the standard evidence $\{X_{-j} = x_{-j}, X_j = y\}$. For each y there is a distinct set of standard forward and backward quantities F_i^y, B_i^y ; however it is easy to see that $F_i^y \equiv F_i$ for $i \leq j-1$ and $B_i^y \equiv B_i$ for $i \geq j$. It follows that $\sum_y \mathbb{P}(S_j = s, \mathcal{E}_y) =$

$$\sum_{y} F_j^y(s) B_j^y(s) = B_j(s) \sum_{y} F_j^y(s) =$$
$$B_j(s) \sum_{y} \sum_{r} F_{j-1}^y(r) \alpha(r,s) \beta(s,y) =$$
$$B_j(s) \sum_{r} F_{j-1}(r) \alpha(r,s).$$

Our main result is a straightforward consequence of the two lemmas above:

Theorem 3: For an arbitrary fixed $j \in \{1, \ldots, n\}$:

$$K_{j} = \sum_{s} \frac{F_{j}^{*}(s)B_{j}(s)}{\sum_{r}F_{j}^{*}(r)B_{j}(r)} \log\left(\frac{F_{j}^{*}(s)}{F_{j}(s)} \cdot \frac{\sum_{r}F_{j}(r)B_{j}(r)}{\sum_{r}F_{j}^{*}(r)B_{j}(r)}\right),$$

where the quantities $(F_i)_{i=1,...,n}, (B_i)_{i=1,...,n}$ and $(F_i^*)_{i=1,...,n}$ are computed once and for all independently of j using the recursions (1), (2), (3). The complexity of computing $(K_j)_{j=1,...,n}$ is $\mathcal{O}(nm^2)$.



Fig. 2: KLD function $(K_j)_{j=1,...,n}$ for the temperature change time series. The five highest K_j are: $K_{1917} = 2.96$, $K_{1915} = 2.30$, $K_{1900} = 1.82$, $K_{1898} = 1.47$, $K_{1914} = 1.46$; the corresponding datapoints are depicted as empty dots.

III. APPLICATION TO TIME SERIES SEGMENTATION

We illustrate the practical interest of our influence measure on a real dataset, namely a time series consisting of 106 annual changes in global temperature between 1880 and 1985 [14]. Following the approach suggested by [9], the dataset can be modeled with an homoscedastic HMM in which each observation follows a Gaussian distribution whose mean depends on the corresponding hidden state; we assume that there are three hidden states. We estimated the parameters of the HMM with the EM algorithm and obtained for the three hidden Gaussian distributions the means $\mu_1 = -0.372$, $\mu_2 = 0.069$, $\mu_3 = -0.068$, and standard deviation $\sigma = 0.114$; moreover the transition matrix is $\pi(i, j) = \eta/2$ if $i \neq j$ and $\pi(i, i) = 1 - \eta$, where the estimated transition rate is $\eta = 0.085$.

Fig. 2 shows the temperature time series together with the KLD K_j for each j. Five years clearly appear to have a greater influence on the posterior distribution of the hidden states: 1917, 1915, 1900, 1898, 1914. It might be interesting to investigate the reasons why these five years are so influential, looking for either specific climatic events or possible changes in the data collection protocol.

In order to validate the findings in Fig. 2, we further investigated the effect of the five most influential observations on the posterior distribution of the segmentation by comparing the marginal posterior distributions obtained with all the observations and after removing the five most influential ones, see Fig. 3. Unsurprisingly, the most dramatic changes occur in the neighborhood of the removed data. When all the observations are taken into account, the period 1880-1920 is characterized by a long segment of negative annual temperature change interrupted by two short periods of slightly positive annual change around years 1900 and 1914 (Fig. 3, top). When the most influential observations are not considered, these two interruptions basically disappear (Fig. 3, bottom).

The KLD measure of influence is hence clearly effective in pointing out observations that have a dramatic effect on the posterior segmentation. These observations can be interpreted either as critical and particularly meaningful data or as outliers (i.e. observations that are not generated by the underlying statistical model).



Fig. 3: Marginal posterior distributions $\mathbb{P}(S_j|\text{obs})$ of the threelevel segmentation considering all the observations (top) and after removing the five most influential ones (bottom). Solid black line: $\mathbb{P}(S_j = 1|\text{obs})$, i.e. x_j is Gaussian with mean $\mu_1 =$ -0.372; dashed red: $\mathbb{P}(S_j = 2|\text{obs})$ with $\mu_2 = 0.069$; dotted blue: $\mathbb{P}(S_j = 3|\text{obs})$ with $\mu_3 = -0.068$.

Application to Outlier Detection

Following [15] (in the context of linear regression), we argue that the KLD-based measure of influence of an observation can be also used for effective outlier detection in data modeled with the HMM. Indeed, if $X_j = x_j$ is an outlier, then it must have a strong influence on the posterior distribution of the hidden variables, which in turn, must differ significantly from the posterior distribution of the hidden variables conditioned on all the observations but X_j . In other words, we expect the KLD distance K_j to be significantly larger when $X_j = x_j$ is an outlier (an illustration supporting this assumption can be found in the Supplementary Material).

In order to explore whether the KLD is an appropriate measure for outlier detection, we considered semi-parametric simulations based on the time series of changes in global temperature described above. The original data is assumed to be free of outliers. 1000 simulations under the null hypothesis H0 (no outliers) were obtained by random sampling n = 106/2 = 53 data points in the original time series. 1000 simulations under the alternative hypothesis H1 (presence of outliers) were obtained by sampling n = 53 data points from the original time series and adding a Gaussian noise $\mathcal{N}(0, \delta^2)$ to each of them with probability 0.05. Hence, the resulting average number of outliers in each H1 simulation is $0.05 \cdot n = 2.65$.

For each simulation, we computed the following global statistics for outlier detection: the maximum K_j , the maximum absolute normalized z-score (using a three component mixture model) and the maximum Local Outlier Factor (LOF) score computed with the R package Rlof [16] after rescaling both

TABLE I: Performance of methods for outlier detection: empirical AUC with 95% confidence intervals, sample size = 400. δ is the standard deviation of the Gaussian noise characterizing outliers.

Method	$\delta = 0.5$	$\delta = 2.0$	$\delta = 3.0$
KLD	0.62[0.57, 0.68]	$0.79\left[0.74, 0.84 ight]$	$0.86\left[0.82, 0.90 ight]$
Z-value	0.58[0.52, 0.64]	$0.61\left[0.55, 0.66 ight]$	$0.59\left[0.53, 0.65 ight]$
LOF	0.73[0.68, 0.78]	0.93[0.90, 0.96]	0.94[0.91,0.96]

year and temperature axes. Details on the statistics can be found in the Supplementary Material.

The performances of the three global statistics for three different values of δ were assessed with the empirical AUC (computed with [17]); the results are depicted in Table I. The statistics based on the Z-value have very poor performance, whereas the KLD-based statistics has a good discriminating power for $\delta \geq 2.0$. However, the method consisting in computing the LOF scores after normalizing both axes prove to be very performant for each value of δ . All three methods are very fast: it takes less than 0.5 seconds for generating a simulation and computing all three statistics.

IV. CONCLUSIONS

An interesting question in Hidden Markov Models is assessing the relative importance of each observation with respect to the sequence of hidden states. In order to measure how influential is the *j*-th observation, we suggest to use the Kullback-Leibler distance K_j between the conditional distribution of the hidden sequence given the whole observation sequence and the conditional distribution of the hidden sequence given all the observations but the *j*-th one. The suggested measure of influence focuses on the posterior distribution of the hidden sequence rather than on the parameter estimate (like in sensitivity analysis) and it is therefore suitable for problems where the information of interest is the hidden sequence (speech recognition [3], genetics [11], bioinformatics [10])

The most important contribution of this letter is a novel linear complexity algorithm for computing the measures of influence of all the observations. Our algorithm is based on simple recursions derived from the forward-backward algorithm for HMMs and can be easily extended in order to take into account pairs, triplets or h consecutive observations. In this case the complexity is $O(nhm^2)$. The algorithm can be also extended to more complex configurations of observations, the resulting complexity depends on the combinatorics of the configuration.

We showed that the KLD influence measure can help to detect outliers in time series modeled by HMMs, the intuition being that anomalies must be more influential than other observations. In this context, the KLD-based method proves to be efficient for global detection, even though it is less performant than specific methods such as the LOF algorithm (after appropriate rescaling).

However, the main interest of the KLD measure of influence is the detection of individual observations which, rather than being outliers, are meaningful values playing a critical role in the problem under consideration. New knowledge can be uncovered by investigating the most influential observations found with our influence measure. For example, in the context of protein structure analysis, structural alphabet are encoded through HMMs [10]. Pointing out highly influential residuals in the encoding through the KLD measure might reveal interesting structural properties (e.g. alternative 3D-structures).

REFERENCES

- Y. Ephraim and N. Merhav, "Hidden Markov processes," *IEEE Transactions on Information Theory*, vol. 48, no. 6, pp. 1518–1569, 2002.
- [2] L. Rabiner, "A tutorial on hidden Markov models and selected applications in speech recognition," *Proceedings of the IEEE*, vol. 77, no. 2, pp. 257–286, 1989.
- [3] B. Gold, N. Morgan, and D. Ellis, Speech and audio signal processing. Wiley Online Library, 2011.
- [4] R. Durbin, S. R. Eddy, A. Krogh, and G. Mitchison, *Biological Sequence Analysis : Probabilistic Models of Proteins and Nucleic Acids*. Cambridge University Press, Jul. 1999.
- [5] C. M. Bishop, Pattern Recognition and Machine Learning (Information Science and Statistics). Secaucus, NJ, USA: Springer-Verlag New York, Inc., 2006.
- [6] M. Do, "Fast approximation of Kullback-Leibler distance for dependence trees and hidden Markov models," *IEEE Signal Processing Letters*, vol. 10, no. 4, pp. 115–118, 2003.
- [7] D. Cook, "Detection of Influential Observations in Linear Inference," Journal of Statistical Planning and Inference, vol. 37, pp. 51–68, 1977.
- [8] W. Johnson, "Influence measures for logistic regression: Another point of view," *Biometrika*, vol. 72, no. 1, pp. 59–65, 1985.
- [9] J. Fridlyand, A. Snijders, D. Pinkel, D. Albertson, and A. Jain, "Hidden Markov models approach to the analysis of array CGH data," *Journal* of *Multivariate Analysis*, vol. 90, no. 1, pp. 132–153, 2004.
- [10] A. Camproux, R. Gautier, P. Tuffery *et al.*, "A hidden markov model derived structural alphabet for proteins," *Journal of molecular biology*, vol. 339, no. 3, pp. 591–606, 2004.
- [11] Y. Li, C. Willer, J. Ding, P. Scheet, and G. Abecasis, "MaCH: using sequence and genotype data to estimate haplotypes and unobserved genotypes," *Genetic epidemiology*, vol. 34, no. 8, pp. 816–834, 2010.
- [12] J. Silva and S. Narayanan, "Upper Bound Kullback-Leibler Divergence for Transient Hidden Markov Models," *IEEE Transactions on Signal Processing*, vol. 56, no. 9, pp. 4176–4188, 2008.
- [13] S.M.E. Sahraeian and B.J. Yoon, "A novel low-complexity HMM similarity measure," *IEEE Signal Processing Letters*, vol. 18, no. 2, pp. 87–90, 2011.
- [14] R. J. Hyndman, "Time Series Data Library," http://data.is/TSDLdemo, Accessed on September, 24 2012.
- [15] S. Chatterjee and A. Hadi, "Influential observations, high leverage points, and outliers in linear regression," *Statistical Science*, vol. 1, no. 3, pp. 379–393, 1986.
- [16] Y. Hu, W. Murray, and Y. Shan, *Rlof: R parallel implementation of Local Outlier Factor (LOF)*, 2011, R package version 1.0.0. [Online]. Available: http://CRAN.R-project.org/package=Rlof
- [17] X. Robin, N. Turck, A. Hainard, N. Tiberti, F. Lisacek, J.-C. Sanchez, and M. Müller, "pROC: an open-source package for R and S+ to analyze and compare ROC curves," *BMC Bioinformatics*, vol. 12, p. 77, 2011.
- [18] M. Markou and S. Singh, "Novelty detection: a review, Part 1: statistical approaches," *Signal Processing*, vol. 83, no. 12, pp. 2481–2497, 2003.
- [19] D. Yeung and Y. Ding, "Host-based intrusion detection using dynamic and static behavioral models," *Pattern recognition*, vol. 36, no. 1, pp. 229–243, 2003.
- [20] D. Zhang, D. Gatica-Perez, S. Bengio, and I. McCowan, "Semisupervised adapted HMMs for unusual event detection," in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition CVPR 2005*, vol. 1. IEEE, 2005, pp. 611–618.
- [21] S. Shah, X. Xuan, R. DeLeeuw, M. Khojasteh, W. Lam, R. Ng, and K. Murphy, "Integrating copy number polymorphisms into array CGH analysis using a robust HMM," *Bioinformatics*, vol. 22, no. 14, pp. e431– e439, 2006.
- [22] M. Siu and A. Chan, "A robust Viterbi algorithm against impulsive noise with application to speech recognition," *IEEE Transactions on Audio*, *Speech, and Language Processing*, vol. 14, no. 6, pp. 2122–2133, 2006.

- [23] S. Chatzis and T. Varvarigou, "A Robust to Outliers Hidden Markov Model with Application in Text-Dependent Speaker Identification," in *IEEE International Conference on Signal Processing and Communications ICSPC 2007.* IEEE, 2007, pp. 804–807.
- [24] S. Chatzis, D. Kosmopoulos, and T. Varvarigou, "Robust sequential data modeling using an outlier tolerant hidden Markov model," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 31, no. 9, pp. 1657–1669, 2009.
- [25] M. Breunig, H. Kriegel, R. Ng, and J. Sander, "LOF: identifying densitybased local outliers," in ACM Sigmod Record, vol. 29, no. 2. ACM, 2000, pp. 93–104.

Vittorio Perduca received the M.Sc. and Ph.D. in mathematics from the University of Turin (Italy) and the M.Sc. in computational biology from Paris Descartes University (France) in 2004, 2009 and 2011 respectively.

Dr. Perduca's postdoctoral fellowship is supported by the *Fondation Sciences Mathématiques de Paris*; his research interests include belief propagation algorithms for Bayesian networks with applications in computational biology.

Gregory Nuel received the Ph.D. in mathematics and the *Habilitation* à *Diriger des Recherches* in 2001 and 2007 respectively, both from the University of Evry (France).

Dr. Nuel currently works as a senior researcher for CNRS. His topics of interests include models with incomplete data, Bayesian networks, motifs in random sequences, and a wide range of biomedical applications.

Appendix

SUPPLEMENTARY MATERIAL WITH TECHNICAL DETAILS

APPLICATION TO OUTLIER DETECTION

In this section, we give more details on our application of the KLD measure of influence to the detection of outliers in HMMs.

A few methods based on the HMM have been developed for outlier detection [18], [19], [20]. In the main paper we consider a related yet different problem, namely the detection of outliers in data *modeled* with the HMM, for instance time series. This problem is not new in the literature, for instance an *ad hoc* model for outliers in data modeled by HMMs was introduced in a Bayesian framework in [21]. Other authors suggested to tackle the problem by means of a robust Viterbi algorithm performing a joint decoding and outlier detection during the Viterbi search [22]. Following [15] (in the context of linear regression) we suggest to detect outliers in HMMs by means of our KLD-based measure of observation influence.

An outlier is an observation that is not generated by the underlying statistical model. Since HMMs are intrinsically heterogeneous, the detection of outliers in data modeled by HMMs is a challenging problem. For instance, change pointdetection methods based on HMMs are known to be particularly sensitive to the presence of outliers in the sense that a single outlier can result in a segment consisting in just one point [23], [24].

As explained in the main paper, we expect the KLD distance K_i to be significantly larger when $X_i = x_i$ is an outlier.

Fig. 4: KLD function $(K_j)_{j=1,...,n}$ for the original time series (top) and for the times series with two outliers artificially added (triangles): $x_{1884} = 0.2$, $x_{1939} = -0.6$.

A. Illustration

Consider the same time series consisting of 106 annual changes in global temperature between 1880 and 1985 as in the main paper. The dataset can be modeled with an HMM in which each observation follows a Gaussian distribution whose mean depends on the corresponding hidden state. The upper plot in Fig. 4 shows the KLD function computed after estimating the parameters with the EM algorithm (same function as in the upper plot of Fig. 2 in the main text but on a different scale). We assume that the original data is outlier free (as we explain in the main text, the peaks in the figure can be interpreted as pointers to meaningful observations). However, if the dataset contains outliers, can we detect them with our measure of influence based on the KLD? In order to answer this question, we manually added two outliers and re-computed the KLD function, after re-estimating the parameters. The results are depicted in the lower plot of Fig. 4 and clearly show that the KLD function has two peaks in the two outliers.

B. Comparison with other methods

We give here the details of the empirical comparison study whose results are reported in the main paper.

Data. We considered semi-parametric simulations based on the time series of changes in global temperature described above. The original data is assumed to be free of outliers. 1000 simulations under the null hypothesis H0 (no outliers) were obtained by random sampling n = 106/2 = 53 data points in the original time series. 1000 simulations under the alternative hypothesis H1 (presence of outliers) were obtained by sampling n = 53 data points from the original time series and adding a Gaussian noise $\mathcal{N}(0, \delta^2)$ to each of them with probability 0.05. Hence, the resulting average number of outliers in each H1 simulation is $0.05 \cdot n = 2.65$.



6

We tested the global hypothesis H1 that the data contain *at least one outlier* against the hypothesis H0 that there are no outliers with the following alternative methods:

KLD-based method. For each simulation q we estimated the parameters in the HMM modeling the dataset with the EM algorithm and then computed the global statistics

$$T_q = \max_{j=1,\dots,n} K_j$$

Z-value. For each simulation q we clustered the data with the k-means algorithm (k = 3) and then computed the Z-value $Z_j = \frac{x_j - \mu}{\sigma}$ of each data point x_j with respect to the mean μ and standard deviation σ of its cluster. We considered the global statistics

$$S_q = \max_{j=1,\dots,n} |Z_j|.$$

Local Outlier Factor (LOF). The LOF algorithm is a density based method [25]. For each data point, the LOF score is calculated by comparing the local density of the point (defined as the inverse of the average distance from its r-nearest neighbors) to the average of the densities of its neighbors. The score is interpreted as a measure of whether the point is in a denser or sparser region of the dataset. A ranking of the points as outliers is obtained by sorting them according to their LOF scores.

The LOF score depends on the choice of the distance parameter r; as suggested in [25] for each point we took the maximal LOF score on a range of integer values for r, namely $r \in \{10, ..., 20\}$. We considered the global statistics

$$L_q = \max_{j=1,\dots,n} \max_{r=10,\dots,20} LOF_r(\tilde{x}_j, \tilde{t}_j),$$

where \tilde{x}_j and \tilde{t}_j are the standardized values of x_j and t_j (i.e. we rescaled both axes before computing the LOF scores). The LOF scores were computed using the R package Rlof [16].

ROC analysis. We assessed the performance of each method by means of the empirical Area Under the Curve (AUC). The AUC measures the surface under the Receiver Operating Characteristic (ROC) curve and can be qualitatively interpreted as follows: AUC ≤ 0.6 means "fail"; 0.6 < AUC ≤ 0.70 means "poor"; 0.7 < AUC ≤ 0.80 means "fair"; 0.8 < AUC ≤ 0.9 means "good"; 0.9 < AUC ≤ 1.0 means "excellent". AUC computations were performed with the R package pROC [17] using the statistics computed for each method and simulation.