

Non-negative Dynamical System with Application to Speech and Audio

Fevotte, C.; Le Roux, J.; Hershey, J.R.

TR2013-021 May 2013

Abstract

Non-negative data arise in a variety of important signal processing domains, such as power spectra of signals, pixels in images, and count data. This paper introduces a novel non-negative dynamical system (NDS) for sequences of such data, and describes its application to modeling speech and audio power spectra. The NDS model can be interpreted both as an adaptation of linear dynamical systems (LDS) to non-negative data, and as an extension of non-negative matrix factorization (NMF) to support Markovian dynamics. Learning and inference algorithms were derived and experiments on speech enhancement were conducted by training sparse non-negative dynamical systems on speech data and adapting a noise model to the unknown noise condition. Results show that the model can capture the dynamics of speech in a useful way.

IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)

This work may not be copied or reproduced in whole or in part for any commercial purpose. Permission to copy in whole or in part without payment of fee is granted for nonprofit educational and research purposes provided that all such whole or partial copies include the following: a notice that such copying is by permission of Mitsubishi Electric Research Laboratories, Inc.; an acknowledgment of the authors and individual contributions to the work; and all applicable portions of the copyright notice. Copying, reproduction, or republishing for any other purpose shall require a license with payment of fee to Mitsubishi Electric Research Laboratories, Inc. All rights reserved.

NON-NEGATIVE DYNAMICAL SYSTEM WITH APPLICATION TO SPEECH AND AUDIO

Cédric Févotte^{*1}, Jonathan Le Roux², John R. Hershey²

¹Laboratoire Lagrange (CNRS, OCA & University of Nice), Parc Valrose, 06000 Nice, France

²Mitsubishi Electric Research Laboratories (MERL), 201 Broadway, Cambridge, MA 02139, USA

ABSTRACT

Non-negative data arise in a variety of important signal processing domains, such as power spectra of signals, pixels in images, and count data. This paper introduces a novel non-negative dynamical system (NDS) for sequences of such data, and describes its application to modeling speech and audio power spectra. The NDS model can be interpreted both as an adaptation of linear dynamical systems (LDS) to non-negative data, and as an extension of non-negative matrix factorization (NMF) to support Markovian dynamics. Learning and inference algorithms were derived and experiments on speech enhancement were conducted by training sparse non-negative dynamical systems on speech data and adapting a noise model to the unknown noise condition. Results show that the model can capture the dynamics of speech in a useful way.

Index Terms— non-negative dynamical system (NDS), linear dynamical system (LDS), multiplicative innovations, non-negative matrix factorization (NMF), source separation.

1. INTRODUCTION

Non-negative sequences arise in a variety of important signal processing domains, such as power spectra of signals, pixels in images, and sequences of count data. This paper introduces a novel dynamical system for non-negative data, and describes its application to speech and audio modeling.

This work bridges two active fields, dynamical systems and non-negative matrix factorization (NMF). Dynamical systems are a long-standing area of research with applications in many scientific fields. A large body of literature is devoted to the case of linear dynamical systems (LDS), which describes an observed sequence $v_n \in \mathbb{R}^F$, indexed by $n \in [1, \dots, N]$, via latent variables $h_n \in \mathbb{R}^K$, according to the equations:

$$h_n = Ah_{n-1} + \xi_n, \quad (1)$$

$$v_n = Wh_n + \epsilon_n, \quad (2)$$

where A and W are matrices of dimensions $K \times K$ and $F \times K$, respectively, and ξ_n and ϵ_n are independent, additive, and typically Gaussian random variables. (1) describes the dynamics of the state variable h_n and (2) describes the observation model. A , W , ξ_n and ϵ_n are referred to as state-transition matrix, dictionary, state innovation and data innovation, respectively. The LDS model does not naturally apply to the case where h_n and v_n are *non-negative*. This paper develops an analogous model for the non-negative domain.

NMF, in contrast, focuses on the case of non-negative data, and has attracted a lot of attention in signal processing and machine learning communities since the publication of the seminal paper [1].

In the general case, NMF is the problem of finding an approximation of the form $V \approx WH$ where V and H are non-negative matrices. This approximation is generally obtained by minimizing a cost function $D(V|WH)$ that measures the dissimilarity between V and WH . Some variants of NMF make additional structure assumptions on either W or H , in the form of regularization terms added to the cost function, that favor properties such as sparsity or smoothness. In some settings, the columns of V form a sequence $v_1 \dots v_N$, and may exhibit evolving dynamics, in the sense of statistical dependencies between elements in the sequence. Standard forms of NMF, which model each element of the sequence independently, as $v_n \approx Wh_n$, will fail to capture such dynamics. From this perspective our work brings probabilistic dynamics to NMF, comparable to that of the traditional linear dynamical system.

The discrete-state Hidden Markov Model (HMM) is another dynamical system that has been commonly used to handle dynamics of speech, for example in automatic speech recognition as well as in speech separation [2]. In this setting, the speech features are usually taken to be cepstral coefficients or other log-spectrum-based features. However, HMMs lead to combinatorial complexity due to the discrete state space, especially in the co-occurrence of several speakers. Because of the discreteness of the state-space and the state-conditional independence of adjacent frames, HMMs also famously do not easily handle gain adaptation and continuity over time. In contrast, standard NMF solves both the computational cost (of linear complexity per iteration) and gain adaptation problems (through H), but it does not handle continuous dynamics. By bringing continuous dynamics to an NMF-like formulation, we hope to obtain the best of both worlds.

The proposed non-negative dynamical system (NDS) model is given by

$$h_n = Ah_{n-1} \circ \xi_n, \quad (3)$$

$$v_n = Wh_n \circ \epsilon_n, \quad (4)$$

where all variables are non-negative with the same dimensions as above, and ‘ \circ ’ denotes element-wise multiplication. The multiplicative innovations ξ_n and ϵ_n are non-negative random variables. The observation model (4) operates similarly to standard NMF, whereas the latent dynamics (3) capture statistical dependency between frames similarly to LDS.

In Section 2, we complete the presentation of the proposed model (3)-(4) with additional statistical assumptions. In Section 3, we present a majorization-minimization (MM) algorithm for maximum a posteriori (MAP) estimation of W , H , A given V . Section 4 discusses how the proposed NDS applies to the modeling of speech spectra and reports speech enhancement results.

^{*}The research presented in this paper was conducted while C. Févotte was a visiting scientist at MERL over Summer 2012.

2. MODEL

2.1. Statistical assumptions

We assume the state and data innovations to be independent and identically distributed (i.i.d), spatially independent and Gamma distributed, such that $p(\xi_1, \dots, \xi_N) = \prod_{k,n} G(\xi_{kn}|\alpha_k, \beta_k)$ and $p(\epsilon_1, \dots, \epsilon_N) = \prod_{f,n} G(\epsilon_{fn}|\nu_f, \delta_f)$, where α_k, β_k, ν_f and δ_f are positive scalars and G refers to the Gamma distribution, with probability density function (pdf) $G(x|\alpha, \beta) = \beta^\alpha/\Gamma(\alpha) x^{\alpha-1}e^{-\beta x}$. It follows that v_n and h_n are conditionally Gamma distributed, such that $p(h_n|Ah_{n-1}) = \prod_k G(h_{kn}|\alpha_k, \beta_k/\sum_j a_{kj}h_{j(n-1)})$ and $p(v_n|Wh_n) = \prod_f G(v_{fn}|\nu_f, \delta_f/\sum_k w_{fk}h_{kn})$. The former holds for $n > 1$ only and we assume independent scale-invariant Jeffreys prior for $n = 1$, i.e., $p(h_{k1}) \propto 1/h_{k1}$. The expectation of the state variables and data under the model are

$$\mathbb{E}(h_{kn}|Ah_{n-1}) = \frac{\alpha_k}{\beta_k} \sum_j a_{kj}h_{j(n-1)}, \quad (5)$$

$$\mathbb{E}(v_{fn}|Wh_n) = \frac{\nu_f}{\delta_f} \sum_k w_{fk}h_{kn}. \quad (6)$$

For simplicity, we will assume $\nu_f = \delta_f = \delta$, so that $\mathbb{E}(V) = WH$, which is a natural assumption underlying most NMF settings (see for example [3]). Under this assumption, the negative log-likelihood $-\log p(V|WH)$ is essentially the Itakura-Saito (IS) divergence from V to WH , as explained in [4]. The expression of $p(h_n|Ah_{n-1})$ reveals a scale ambiguity between β_k and $\{a_{kj}\}_j$, which we fix by setting $\beta_k = \alpha_k$, implying that $\mathbb{E}(h_n|Ah_{n-1}) = Ah_{n-1}$.

The multiplicative and non-negative state and data innovations preserve the non-negativity in the generative model. Additive real-valued Gaussian innovations, as used in LDS, would fail to do so. Additive non-negative innovations would ensure non-negativity, but would only allow monotonically increasing values. In contrast, for multiplicative innovations that can take values both lower and greater than one, the coefficients of v_n and h_n are allowed to increase and decrease.

2.2. Related work

To the best of our knowledge there is no prior work about the NDS (3)-(4) proposed in this paper. There is some literature on non-negative dynamical systems, where the state-space matrices, state and observation vectors are assumed non-negative. However, to the best of our knowledge, only non-negative additive perturbations have been considered, see, e.g., [5, 6]. Furthermore most of the literature we found on the topic is devoted to the theoretical properties of these systems (stability, observability) rather than inference in or application of these models. The closest to our model is perhaps [7], which studies the conditions for existence, uniqueness and stability of the system defined by Eq. (3) alone, but does not provide algorithms nor procedures for inference.

A special case of the NDS (3)-(4) has been addressed in the NMF literature. In [4, 8], algorithms were derived for the special case $A = I_K$, where I_K denotes the identity matrix of dimension K , and $\delta = 1$ (multiplicative exponentially-distributed observation innovation). In these papers, the state innovation was arbitrarily assumed either Gamma or inverse-Gamma with mode obtained at 1, i.e., the mode of $p(h_{kn}|h_{k(n-1)})$ is obtained at $h_{kn} = h_{k(n-1)}$. The resulting method, coined smooth Itakura-Saito NMF, allows to regularize the individual rows of H temporally, assuming mutual independence of the rows. The proposed model goes an important

step forward by lifting the mutual independence assumption, which is generally not a realistic one. Indeed, the activation of a pattern at frame n is also likely to correlate (or anti-correlate) with the activation of other patterns at frame $n-1$. This is what the proposed model achieves through the introduction of the state-transition matrix A in Eq. (3).

A preliminary attempt at introducing dynamics into NMF has been made by Smaragdis *et al.*, for speech denoising [9] and sound classification [10]. They compute standard NMF decompositions from spectral training data describing each sound class (e.g., in the denoising setting, speech and noise) and compute the average temporal dynamics of the returned matrices H for each class *a posteriori*. Test sounds are then decomposed onto the learned spectral patterns, using a regularization step that employs the precomputed temporal statistics of each class. In particular [10] employs an ad-hoc forward-backward smoothing of the activations. The work presented in this paper pursues a more formal approach, and proposes a well-posed statistical model for non-negative data along with principled algorithms for inference.

3. MAXIMUM A POSTERIORI ESTIMATION

We have derived a majorization-minimization (MM) algorithm for MAP estimation of the parameters W, A, H given user-defined values of the remaining parameters δ and α . In the following, we only present the main steps of the procedure to meet the space limitations constraints. The MAP objective function is defined by $C(W, H, A) = -\log p(V|WH) - \log p(H|A)$. Our MM algorithm is a block-coordinate descent algorithm that updates W, H and A individually and conditionally upon the current values of the other parameters. The algorithm alternates between forming an upper bound of the objective function at the current parameter settings, and optimizing parameters to minimize this bound.

Let \tilde{W}, \tilde{H} and \tilde{A} denote the parameter values at the current iteration and consider for example the update of W . Denote for example $F(W) = C(W, \tilde{H}, \tilde{A})$ the function to be minimized w.r.t W . The first step of the MM algorithm consists in building an upper bound $G(W, \tilde{W})$ of $F(W)$ which is tight for $W = \tilde{W}$, i.e., $F(W) \leq G(W, \tilde{W})$ for all W and $G(\tilde{W}, \tilde{W}) = F(\tilde{W})$. The second step consists in minimizing the bound w.r.t W , producing a valid descent algorithm. Indeed, at iteration $i+1$, it holds by construction that $F(W^{(i+1)}) \leq G(W^{(i+1)}, W^{(i)}) \leq G(W^{(i)}, W^{(i)}) = F(W^{(i)})$. The same principle applies to the updates of H and A .

The upper bounds can be derived using standard inequalities, namely Jensen's inequality for the convex parts of the objective functions and the tangent inequality for the concave parts. Using this strategy, the updates of W and A are multiplicative and given by

$$w_{fk} = \tilde{w}_{fk} \sqrt{\frac{\sum_{n=1}^N h_{kn} v_{fn} / \tilde{v}_{fn}^2}{\sum_{n=1}^N h_{kn} / \tilde{v}_{fn} + \lambda}} \quad (7)$$

$$a_{kj} = \tilde{a}_{kj} \sqrt{\frac{\beta_i \sum_{n=2}^N h_{j(n-1)} h_{kn} / \tilde{g}_{kn}^2}{\alpha_i \sum_{n=2}^N h_{j(n-1)} / \tilde{g}_{kn}}} \quad (8)$$

where $\tilde{v}_{fn} = \sum_k \tilde{w}_{fk} h_{kn}$, $\tilde{g}_{kn} = \sum_j \tilde{a}_{kj} h_{j(n-1)}$ and λ is a constant that prevents degenerate solutions such that $\|W\| \rightarrow \infty$ and $\|H\| \rightarrow 0$. Owing to the Markovian structure of H , adjacent columns are coupled in the optimization. We have employed a left-to-right block-coordinate descent approach that updates h_n at iteration i conditionally on $h_{n-1}^{(i)}$ and $h_{n+1}^{(i-1)}$, for $1 < n < N$. With this

approach, updates of h_{kn} are available in closed form and merely involve rooting a polynomial of order 2, such that

$$h_{kn} = \frac{\sqrt{b_{kn}^2 - 4a_{kn}c_{kn}} - b_{kn}}{2a_{kn}}, \quad (9)$$

where $a_{kn} = \delta \sum_f \frac{w_{fk}}{\tilde{v}_{fn}} + \sum_j \alpha_j \frac{a_{jk}}{\tilde{g}_{j(n+1)}} + \frac{\beta_k}{g_{kn}}$, $b_{kn} = 1 - \alpha_k$, $c_{kn} = -\tilde{h}_{kn}^2 \left(\delta \sum_f w_{fk} \frac{v_{fn}}{\tilde{v}_{fn}^2} + \sum_j \beta_j \frac{a_{jk} h_{j(n+1)}}{\tilde{g}_{j(n+1)}^2} \right)$, and $\tilde{v}_{fn} = \sum_k w_{fk} \tilde{h}_{kn}$, $\tilde{g}_{k(n+1)} = \sum_j a_{kj} \tilde{h}_{jn}$, $g_{kn} = \sum_j a_{kj} h_{j(n-1)}$.

4. APPLICATION TO SPEECH

4.1. Spectral modeling of speech with NDS

When $\delta = 1$, such that ξ_{kn} is exponentially distributed, Eq. (4) can be related to a generative model of the power spectrogram in the following Gaussian composite model (GCM). Let x_{fn} denote the complex-valued short-time Fourier transform (STFT) of some time domain audio signal, where f is a frequency bin index and n indexes time frames. The GCM is defined by $x_{fn} = \sum_k c_{fkn}$ and $c_{fkn} \sim N_c(0, w_{fk} h_{kn})$, where $N_c(0, \lambda)$ refers to the circular complex Gaussian distribution with zero mean. The latent components $\{c_{fkn}\}$ can trivially be marginalized from the generative model, yielding $x_{fn} \sim N_c(0, \sum_k w_{fk} h_{kn})$. It follows that the power spectrogram $v_{fn} = |x_{fn}|^2$ of x_{fn} is exponentially distributed with mean $\sum_k w_{fk} h_{kn}$, and can thus be written as Eq. (4) with $\delta = 1$. Note that, when necessary, minimum mean squares estimate (MMSE) of the latent components can be obtained by Wiener filtering and given by

$$\hat{c}_{fkn} = \frac{w_{fk} h_{kn}}{\sum_j w_{fj} h_{jn}} x_{fn}. \quad (10)$$

The GCM has found successful applications in audio source separation and music transcription [4, 11, 12, 13, 14], and generalizes earlier two-component models used in spectral-based audio denoising, e.g., [15]. In contrast with Gaussian mixture models (GMMs) or HMMs, which are prevalent in speech log-spectral modeling and where each data frame is assumed to be in one among many possible states each characterized by a given covariance, the GCM assumes that each data frame is a sum of zero-mean Gaussian-distributed components. In this paper the GCM is used as an observation process for the NDS model.

4.2. Speech enhancement with NDS

We consider a speech enhancement scenario where the time-domain data x_t is a clean speech signal s_t corrupted by additive noise b_t , such that

$$x_t = s_t + b_t, \quad (11)$$

and we wish to produce a speech estimate \hat{s}_t of s_t . Given a trained NDS model of speech ($W^{\text{train}}, A^{\text{train}}$), representative of the unseen source s , and given the corrupted data x , we estimate the non-negative decomposition

$$V \approx W^{\text{train}} H + W^{\text{noise}} H^{\text{noise}} \quad (12)$$

of the power spectrogram V of the noisy observation x , where $W^{\text{train}} H$ represents the speech spectrogram estimate and $W^{\text{noise}} H^{\text{noise}}$ the noise spectrogram estimate. We then reconstruct the time-domain source estimate \hat{s} by MMSE estimation, which

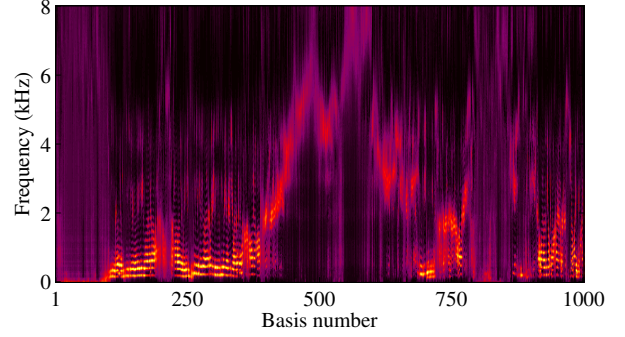


Fig. 1. Example of learned W on female speech (frame length: 512), sorted greedily by similarity.

amounts to Wiener filtering in our model. That is, we take the inverse STFT of

$$\hat{S} = \frac{W^{\text{train}} H}{W^{\text{train}} H + W^{\text{noise}} H^{\text{noise}}} \circ X, \quad (13)$$

where the fraction bar is here elementwise.

In our approach, we assume the source and noise STFTs to follow a GCM, such that $s_{fn} \sim N_c(0, \sum_k w_{fk}^{\text{train}} h_{kn})$ and $b_{fn} \sim N_c(0, \sum_k w_{fk}^{\text{noise}} h_{kn}^{\text{noise}})$. H follows the dynamics of Eq. (3) with transition matrix A^{train} . We recall that $W^{\text{train}}, A^{\text{train}}$ are fixed variables, learned in a training phase. The matrices W^{noise} and H^{noise} are here assumed unknown and with no particular structure, and learned from the data. With these assumptions, the decomposition (12) shall be obtained by minimizing the divergence $D_{IS}(V|W^{\text{train}} H + W^{\text{noise}} H^{\text{noise}})$ penalized by the dynamical term $-\log p(H|A^{\text{train}})$, w.r.t H, W^{noise} and H^{noise} . This can be achieved with an MM algorithm using minor modifications of the derivations presented in Section 3. Besides the use of the IS divergence and the dynamical penalty term, the proposed procedure resembles the semi-supervised NMF setting described in [16].

4.3. Enhancement results

The sampling rate was 16 kHz. Time-frequency analysis was performed using frame lengths of 512, 640, 800 and 960 samples, using for each length a 50% overlap and a sine window for analysis and re-synthesis. For each window length, NDS models were trained separately for male and female speech, each on 1000 utterances (about 50 minutes) from the TIMIT training set. The number of bases was set to $K = 1000$ and the Gamma distribution parameter to $\alpha = 0.01$. This low setting of α strongly encourages sparsity in H , thus encouraging a holistic representation of the spectrum, and a sparse transition matrix. However, sparsity can lead to jitter in the estimated signals, a potential trade-off that will need to be investigated in future work. We show in Figs. 1 and 2 examples of W and A trained on female speech with a frame length of 512.

We evaluated the enhancement algorithm described above on mixtures of 10 speech files by different speakers (5 male and 5 female) from the TIMIT test set with 15 environmental texture sounds from [17] at 3 different input signal-to-noise ratios (SNR), for a total of 450 mixtures. The training and test sets had disjoint sets of speakers. The texture sounds include a wide variety of environmental sounds such as fire, bees, water stream, helicopter, applause, babble noise, shaking paper, etc. The speech test files were 1.8 s to 4.5 s long, and were added in the middle of the 7 s long environmental sounds. The number of noise bases was set to $K^{\text{noise}} = 2$. For each

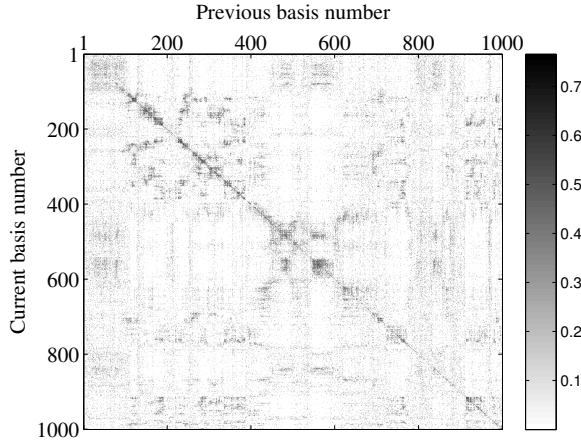


Fig. 2. Example of learned A on female speech (frame length: 512).

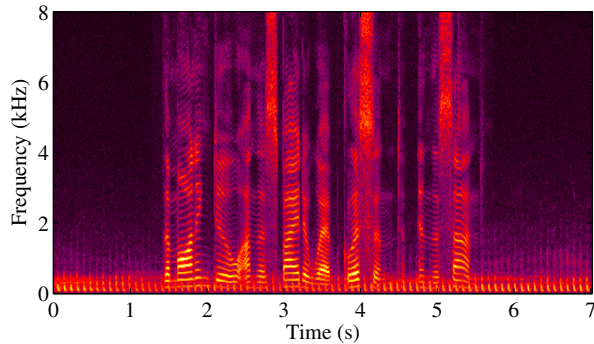


Fig. 3. Mixture of female speech and helicopter sound (10 dB SNR).

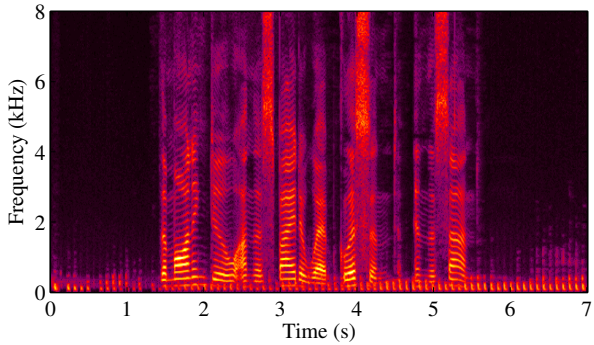


Fig. 4. Result of applying OMLSA on the mixture shown in Fig. 3.

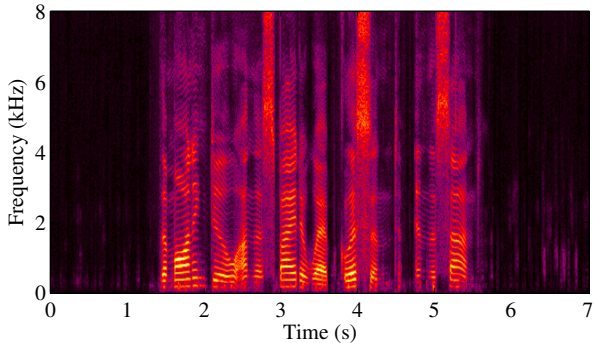


Fig. 5. Result of applying NDS on the mixture shown in Fig. 3.

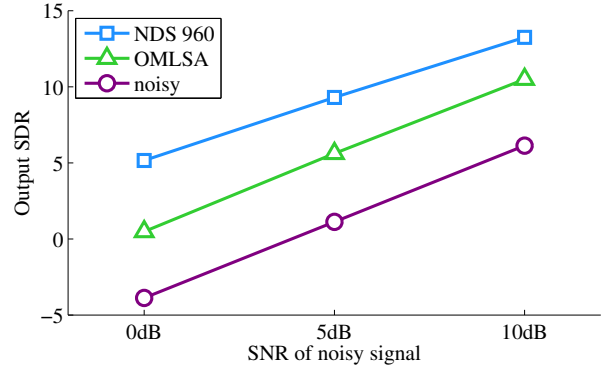


Fig. 6. Signal to distortion ratios for reconstructed speech for NDS with 60 ms windows (NDS 960). For comparison the noisy speech and OMLSA are included.

mixture, we assumed the gender known and used the NDS speech model for that gender.

For comparison, we also show results for the state-of-the-art algorithm combining Optimally-Modified Log Spectral Amplitude Estimator and Improved Minima Controlled Recursive Averaging [18, 19], denoted as OMLSA. To illustrate the behavior of the algorithms, we show in Fig. 3 the spectrogram of a mixture of speech by a female speaker and a helicopter sound at 10 dB SNR, and in Figs. 4 and 5 the output of OMLSA and NDS on that mixture. NDS is able to suppress the non-stationary helicopter noise, while OMLSA fails to do so. The outputs of the algorithms were also quantitatively evaluated using the *bss_eval* toolbox [20], treating our denoising problem as a source separation problem where s is the target source and b is the interfering source, and the perception of speech quality (PESQ) measure [21].

The *bss_eval* results are given in terms of signal-to-distortion ratio (SDR) and shown in Fig. 6. The proposed NDS algorithm with 60 ms windows significantly outperforms OMLSA for all input SNRs. The results for other frame lengths were similar in terms of SDR. When measured by PESQ, in contrast, the scores generally increased with window size, but yielded no significant improvement compared to OMLSA for the window sizes investigated. In informal listening tests, the NDS model was very good at removing a wide range of non-stationary noises, but suffered from a tendency to leave behind residual speech-like sounds. In particular, due to the fact that training data contained lip smacks, breath sounds, etc, noises resembling these were often passed unsuppressed into the speech estimate. We are currently investigating whether these artifacts can be controlled by alternative parameterizations of the model.

5. CONCLUSIONS

We presented a novel non-negative dynamical system called NDS to model sequences of non-negative data, explained its relationships with previous work, derived an efficient MAP estimation algorithm, and explained how it can be applied to speech and audio modeling. In preliminary experiments on a speech enhancement task with real environmental sounds, the proposed NDS algorithm reached similar performance in terms of PESQ scores compared to the state-of-the-art algorithm, and significantly outperformed the state of the art in terms of SDR. Future works include testing more thoroughly the potential of the proposed model on non-negative data of different modalities, developing extensions enabling the sharing of bases across states, and investigating usage of NDS as a noise model.

6. REFERENCES

- [1] D. D. Lee and H. S. Seung, "Learning the parts of objects with nonnegative matrix factorization," *Nature*, vol. 401, pp. 788–791, 1999.
- [2] J. R. Hershey, S. J. Rennie, P. A. Olsen, and T. T. Kristjansson, "Super-human multi-talker speech recognition: A graphical modeling approach," *Computer Speech and Language*, vol. 24, no. 1, pp. 45–66, Jan. 2010.
- [3] V. Y. F. Tan and C. Févotte, "Automatic relevance determination in nonnegative matrix factorization with the beta-divergence," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, in press.
- [4] C. Févotte, N. Bertin, and J.-L. Durrieu, "Nonnegative matrix factorization with the Itakura-Saito divergence. With application to music analysis," *Neural Computation*, vol. 21, no. 3, pp. 793–830, Mar. 2009.
- [5] J. M. van den Hof, "Realization of positive linear systems," *Linear Algebra and its Applications*, vol. 256, no. 0, pp. 287 – 308, 1997.
- [6] W. M. Haddad and V. Chellaboina, "Stability and dissipativity theory for nonnegative dynamical systems: a unified analysis framework for biological and physiological systems," *Nonlinear Analysis: Real World Applications*, vol. 6, no. 1, pp. 35–65, 2005.
- [7] J. Stachurski, "Economic dynamical systems with multiplicative noise," *Journal of Mathematical Economics*, vol. 39, no. 1-2, pp. 135–152, 2003.
- [8] C. Févotte, "Majorization-minimization algorithm for smooth Itakura-Saito nonnegative matrix factorization," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Prague, Czech Republic, May 2011.
- [9] K. W. Wilson, B. Raj, P. Smaragdis, and A. Divakaran, "Speech denoising using nonnegative matrix factorization with priors," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Apr. 2008, pp. 4029–4032.
- [10] J. Nam, G. J. Mysore, and P. Smaragdis, "Sound recognition in mixtures," in *Proc. International Conference on Latent Variable Analysis & Independent Component Analysis (LVA/ICA)*, Tel-Aviv, Israel, 2012.
- [11] S. A. Abdallah and M. D. Plumbley, "Polyphonic transcription by nonnegative sparse coding of power spectra," in *Proc. 5th International Symposium Music Information Retrieval (ISMIR)*, Barcelona, Spain, Oct. 2004, pp. 318–325.
- [12] L. Benaroya, R. Gribonval, and F. Bimbot, "Non negative sparse representation for Wiener based source separation with a single sensor," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, Hong Kong, 2003, pp. 613–616.
- [13] R. M. Parry and I. Essa, "Phase-aware non-negative spectrogram factorization," in *Proc. International Conference on Independent Component Analysis and Signal Separation (ICA)*, London, UK, Sep. 2007, pp. 536–543.
- [14] A. Ozerov and C. Févotte, "Multichannel nonnegative matrix factorization in convolutive mixtures for audio source separation," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 18, no. 3, pp. 550–563, Mar. 2010.
- [15] Y. Ephraim and D. Malah, "Speech enhancement using a minimum-mean square error short-time spectral amplitude estimator," *IEEE Transactions on Acoustics, Speech and Signal Processing*, vol. 32, no. 6, pp. 1109–1121, 1984.
- [16] P. Smaragdis, B. Raj, and M. V. Shashanka, "Supervised and semi-supervised separation of sounds from single-channel mixtures," in *Proc. 7th International Conference on Independent Component Analysis and Signal Separation (ICA)*, London, UK, Sep. 2007.
- [17] J. H. McDermott and E. P. Simoncelli, "Sound texture perception via statistics of the auditory periphery: Evidence from sound synthesis," *Neuron*, vol. 71, pp. 926–940, 2011.
- [18] I. Cohen, "Optimal speech enhancement under signal presence uncertainty using log-spectral amplitude estimator," *IEEE Signal Processing Letters*, vol. 9, no. 4, pp. 113–116, 2002.
- [19] I. Cohen, "Noise spectrum estimation in adverse environments: Improved minima controlled recursive averaging," *IEEE Transactions on Speech and Audio Processing*, vol. 11, no. 5, pp. 466–475, 2003.
- [20] E. Vincent, R. Gribonval, and C. Févotte, "Performance measurement in blind audio source separation," *IEEE Transactions on Audio, Speech and Language Processing*, vol. 14, no. 4, pp. 1462–1469, Jul. 2006.
- [21] A. W. Rix, J. G. Beerends, M. P. Hollier, and A. P. Hekstra, "Perceptual evaluation of speech quality (PESQ)—a new method for speech quality assessment of telephone networks and codecs," in *Proc. IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 2001, vol. 2, pp. 749–752.