

Optimal Estimation of Multivariate ARMA Models

Martha White, Junfeng Wen, Michael Bowling and Dale Schuurmans

Department of Computing Science, University of Alberta, Edmonton AB T6G 2E8, Canada
 {whitem,junfeng.wen,mbowling,daes}@ualberta.ca

Abstract

Autoregressive moving average (ARMA) models are a fundamental tool in time series analysis that offer intuitive modeling capability and efficient predictors. Unfortunately, the lack of globally optimal parameter estimation strategies for these models remains a problem: application studies often adopt the simpler autoregressive model that can be easily estimated by maximizing (a posteriori) likelihood. We develop a (regularized, imputed) maximum likelihood criterion that admits efficient global estimation via structured matrix norm optimization methods. An empirical evaluation demonstrates the benefits of globally optimal parameter estimation over local and moment matching approaches.

Introduction

A central problem in applied data analysis is time series modeling—estimating and forecasting a discrete-time stochastic process—for which the autoregressive moving average (ARMA) and stochastic ARMA (Thiesson et al. 2012) are fundamental models. An ARMA model describes the behavior of a linear dynamical system under latent Gaussian perturbations (Brockwell and Davis 2002; Lütkepohl 2007), which affords intuitive modeling capability, efficient forecasting algorithms, and a close relationship to linear Gaussian state-space models (Katayama 2006, pp.5–6).

Unfortunately, estimating the parameters of an ARMA model from an observed sequence is a computationally difficult problem: no efficient algorithm is known for computing the parameters that maximize the marginal likelihood of the observed data in an ARMA, stochastic ARMA or linear Gaussian state-space model. Consequently, heuristic local estimators are currently deployed in practice (Hannan and Kavalieris 1984; Durbin 1960; Bauer 2005; Lütkepohl 2007; Thiesson et al. 2012), none of which provide a guarantee of how well the globally optimal parameters are approximated. For estimating linear Gaussian state-space models, it has been observed that local maximization of marginal likelihood tends to find local optima that yield poor results (Katayama 2006, Sec. 1.3).

In response to the difficulty of maximizing marginal likelihood, there has been growing interest in *method of mo-*

ments based estimators for state-space models, which offer computationally efficient estimation strategies and sound consistency properties (Andersson 2009; Hsu, Kakade, and Zhang 2012; Anandkumar, Hsu, and Kakade 2012). For ARMA models, the most applicable such estimators are the subspace identification methods for estimating state-space models (Katayama 2006; Moonen and Ramos 1993; Van Overschee and De Moor 1994; Viberg 1995; Song et al. 2010; Boots and Gordon 2012). The statistical efficiency of moment matching, however, generally does not match that of maximum likelihood, which is known to be asymptotically efficient under general conditions (Cramér 1946, Ch. 33). In fact, evidence suggests that the statistical efficiency of current moment matching estimators is quite weak (Foster, Rodu, and Ungar 2012; Zhao and Poupart 2014).

In this paper, we develop a tractable approach to maximum likelihood parameter estimation for stochastic multivariate ARMA models. To efficiently compute a globally optimal estimate, the problem is re-expressed as a regularized loss minimization, which then allows recent algorithmic advances in sparse estimation to be applied (Shah et al. 2012; Candes et al. 2011; Bach, Mairal, and Ponce 2008; Zhang et al. 2011; White et al. 2012). Although there has been recent progress in global estimation for ARMA, such approaches have either been restricted to single-input single-output systems (Shah et al. 2012), estimating covariance matrices for scalar ARMA (Wiesel, Bibi, and Globerson 2013) or using AR to approximate a scalar ARMA model (Anava et al. 2013). By contrast, this paper offers the first efficient maximum likelihood approach to estimating the parameters of a stochastic multivariate ARMA(p, q) model. This convex optimization formulation is general, enabling generalized distributional assumptions and estimation on multivariate data, which has been much less explored than scalar ARMA. An experimental evaluation demonstrates that globally optimal parameters under the proposed criterion yield superior forecasting performance to alternative estimates, including local minimization for ARMA estimation and moment-based estimation methods for state-space models.

Background

An ARMA model is a simple generative model of the form depicted in Figure 1(a), where the innovation variables, $\epsilon_t \in \mathbb{R}^n$, are assumed to be i.i.d. Gaussian, $\mathcal{N}(0, \Sigma)$, and




Figure 1: Graphical models depicting the dependence structure of **(a)** an ARMA(1, 2) model, **(b)** a state-space model. These models are equivalent if the state-space model is in observability canonical form (Benveniste, Metivier, and Priouret 2012, Sec. 6.2.1). Distinct methods are used for estimation depending on whether the variables are discrete or continuous.

the observable variables, $\mathbf{x}_t \in \mathbb{R}^n$, are assumed to be generated by the linear relationship

$$\mathbf{x}_t = \sum_{i=1}^p A^{(i)} \mathbf{x}_{t-i} + \sum_{j=1}^q B^{(j)} \varepsilon_{t-j} + \varepsilon_t. \quad (1)$$

An ARMA(p, q) model is thus parameterized by $A^{(1)}, \dots, A^{(p)} \in \mathbb{R}^{n \times n}$, $B^{(1)}, \dots, B^{(q)} \in \mathbb{R}^{n \times n}$, and a positive semi-definite matrix Σ ; which we simply collect as $\Theta = (\{A^{(i)}\}, \{B^{(j)}\}, \Sigma)$.¹

One classical motivation for ARMA models arises from the Wold representation theorem (Wold 1938), which states that any stationary process can be represented as an infinite sum of innovations plus a deterministic process that is a projection of a current observation onto past observations: $\mathbf{x}_t = p(\mathbf{x}_t | \mathbf{x}_{t-1}, \dots) + \sum_{j=0}^{\infty} B^{(j)} \varepsilon_{t-j}$. Thus the autoregressive component of an ARMA model is often motivated as a more parsimonious representation of this Wold representation (Scargle 1981).

Time series models are used primarily for forecasting: Given an ARMA model with parameters Θ , the value of a future observation \mathbf{x}_{T+h} can be predicted from an observed history $\mathbf{x}_{1:T}$ by evaluating $E[\mathbf{x}_{T+h} | \mathbf{x}_{1:T}, \Theta]$. The key advantage of ARMA is that such forecasts can be computed efficiently; see Appendix ?? for additional details.

Although forecasting is efficient, the problem of estimating the parameters of an ARMA model raises significant computational challenges, which provides the main focus of this paper. To begin, consider the marginal log-likelihood of an observed history $\mathbf{x}_{1:T}$ given a set of parameters Θ :

$$\log p(\mathbf{x}_{1:T} | \Theta) = \sum_{t=1}^T \log p(\mathbf{x}_t | \mathbf{x}_{1:t-1}, \Theta). \quad (2)$$

Despite the fact that the conditional expectation $E[\mathbf{x}_t | \mathbf{x}_{1:t-1}, \Theta]$ can be computed efficiently, the quantity $\log p(\mathbf{x}_t | \mathbf{x}_{1:t-1}, \Theta)$ is not concave in Θ (Mauricio 1995; Lütkepohl 2007, Sec. 12.2-3), which suggests that maximizing the marginal likelihood is a hard computational problem. Another source of difficulty is that ARMA is a latent variable model, hence marginalizing over the unobserved innovations $\varepsilon_{1:T}$ might also be problematic. Given innovations $\varepsilon_{1:T} = [\varepsilon_1, \dots, \varepsilon_T]$, however, $p(\mathbf{x}_t | \mathbf{x}_{1:t-1}, \varepsilon_{1:t-1}, \Theta)$ is a simple Gaussian with mean

$$\boldsymbol{\mu}_t = \sum_{i=1}^p A^{(i)} \mathbf{x}_{t-i} + \sum_{j=1}^q B^{(j)} \varepsilon_{t-j} \quad (3)$$

and covariance Σ . To obtain such a simplified form, we will first characterize the entire data likelihood in terms of the

¹Note that we use the term ARMA to include vector ARMA, with no restriction to scalar time series.

innovations which enables application of the widely used expectation-maximization algorithm.

Lemma 1. *For an auxiliary density $q(\cdot)$ over $\varepsilon_{1:T}$, and entropy $H(q(\cdot))$, it follows that (proof given in Appendix ??):*

$$\begin{aligned} \log p(\mathbf{x}_{1:T} | \Theta) &= \log \int p(\mathbf{x}_{1:T}, \varepsilon_{1:T} | \Theta) d\varepsilon_{1:T} \\ &= \max_{q(\cdot)} \int q(\varepsilon_{1:T}) \log p(\mathbf{x}_{1:T} \varepsilon_{1:T} | \Theta) d\varepsilon_{1:T} + H(q(\cdot)). \end{aligned} \quad (4)$$

The maximum likelihood problem can now be re-expressed as $\min_{\Theta} \min_{\{q(\cdot)\}} -\log p(\mathbf{x}_{1:T} | \Theta)$ where in a standard EM algorithm, the M step would consist of optimizing Θ given $\{q(\cdot)\}$, and the E step would consist of (implicitly) optimizing $\{q(\cdot)\}$ given Θ (Neal and Hinton 1998). A standard variant of the log likelihood in (4) can then be obtained simply by dropping the entropy regularizer $H(q(\cdot))$. This leads to the minimization selecting a Dirac delta distribution on $\varepsilon_{1:T}$ and a far simpler formulation, sometimes known as “hard EM” or “Viterbi EM” (Brown et al. 1993):

$$\begin{aligned} \min_{\Theta} \min_{\varepsilon_{1:T}} -\log p(\mathbf{x}_{1:T}, \varepsilon_{1:T} | \Theta) \\ = \min_{\Theta} \min_{\varepsilon_{1:T}} -\sum_{t=1}^T \left[\log p(\varepsilon_t | \mathbf{x}_{1:t}, \varepsilon_{1:t-1}, \Theta) \right. \\ \left. + \log p(\mathbf{x}_t | \mathbf{x}_{1:t-1}, \varepsilon_{1:t-1}, \Theta) \right]. \end{aligned} \quad (5)$$

This formulation suggests an approach where one successively imputes values $\varepsilon_{1:T}$ for the unobserved innovation variables, then optimizes the parameters. Interestingly, $p(\varepsilon_t | \mathbf{x}_{1:t}, \varepsilon_{1:t-1}, \Theta)$ is a Dirac delta distribution; ε_t must be the residual, $\varepsilon_t = \mathbf{x}_t - \boldsymbol{\mu}_t$, otherwise the loss becomes unbounded. This distribution, therefore, imposes a constraint on the minimization. To maximize likelihood under this constraint, we optimize

$$\begin{aligned} \min_{\Theta} \min_{\varepsilon_{1:T}: \varepsilon_t = \mathbf{x}_t - \boldsymbol{\mu}_t} -\sum_{t=1}^T \log p(\mathbf{x}_t | \mathbf{x}_{1:t-1}, \varepsilon_{1:t-1}, \Theta) \quad (6) \\ = \min_{\Theta, \varepsilon_{1:T}: \varepsilon_t = \mathbf{x}_t - \boldsymbol{\mu}_t} \frac{T}{2} \log((2\pi)^n |\Sigma|) + \frac{1}{2} \sum_{t=1}^T \left\| \Sigma^{-\frac{1}{2}} (\mathbf{x}_t - \boldsymbol{\mu}_t) \right\|^2. \end{aligned} \quad (7)$$

Unfortunately, this optimization raises an important challenge. Due to the direct interaction between Σ , B and $\varepsilon_{1:T}$ the final form of the problem (7) is still not convex in the parameters Θ and $\varepsilon_{1:T}$ jointly. A typical strategy is therefore to first estimate the innovations $\varepsilon_{1:T}$ directly from data, for example by using the errors of a learned autoregressive model, then observing that with the innovation variables fixed and Σ approximated from the innovations, the problem becomes convex in Θ (Hannan and Kavalieris 1984; Lütkepohl 2007). Another more contemporary approach is to convert the ARMA model into a state-space model (see

Figure 1(b)) and then solve for parameters in that model using system identification approaches (Bauer 2005). Though this has been an important advance for efficient ARMA estimation, these approaches still result in local minima.

Regularized ARMA modeling

To develop a likelihood based criterion that admits efficient global optimization, we begin by considering a number of extensions to the ARMA model. First, notice that the ARMA model in (1) can be equivalently formulated by introducing a $B^{(0)}$ and taking $\varepsilon_t \sim \mathcal{N}(\boldsymbol{\mu} = \mathbf{0}, \Sigma = I)$, giving $B^{(0)}\varepsilon_t \sim \mathcal{N}(\mathbf{0}, B^{(0)}B^{(0)'})$. Second, following Thiesson et al. (2012), an independent noise term $\boldsymbol{\eta}_t$ can be added to (1) to obtain the *stochastic ARMA* model

$$\mathbf{x}_t = \sum_{i=1}^p A^{(i)} \mathbf{x}_{t-i} + \sum_{j=0}^q B^{(j)} \varepsilon_{t-j} + \boldsymbol{\eta}_t. \quad (8)$$

A key challenge in estimating the parameters of a classical ARMA model (1) is coping with the deterministic constraint that $\boldsymbol{\eta}_t = \mathbf{0}$, which forces the innovations ε_t to match the residuals (7). The stochastic ARMA model (8) relaxes this assumption by allowing $\boldsymbol{\eta}_t$ to be generated by a smooth exponential family distribution, such as $\boldsymbol{\eta}_t \sim \mathcal{N}(\mathbf{0}, Q_t)$ for covariance Q_t ; smaller Q_t yields a closer approximation to the original ARMA model. Thiesson et al. (2012) have shown that expectation-maximization (EM) updates are only meaningful for non-zero Q_t ; else, EM stops after one iteration. EM is not however guaranteed to find a globally optimal parameter estimate for the stochastic ARMA model. A key advantage of this model, however, is that it allows a convenient re-expression of the marginal log-likelihood (6) by applying the chain rule in the opposite order for \mathbf{x}_t and ε_t :

$$\begin{aligned} (6) &= \min_{\Theta} \min_{\varepsilon_{1:T}} \\ & - \sum_{t=1}^T \left[\log p(\mathbf{x}_t | \mathbf{x}_{1:t-1}, \varepsilon_{1:t}, \Theta) + \log p(\varepsilon_t | \mathbf{x}_{1:t-1}, \varepsilon_{1:t-1}, \Theta) \right] \\ &= \min_{\Theta} \min_{\varepsilon_{1:T}} - \sum_{t=1}^T \left[\log p(\mathbf{x}_t | \mathbf{x}_{1:t-1}, \varepsilon_{1:t}, \Theta) + \log p(\varepsilon_t | \Theta) \right], \end{aligned}$$

since ε_t is independent of past innovations and data without \mathbf{x}_t . Furthermore, $p(\varepsilon_t | \Theta) = p(\varepsilon_t)$ since the covariance was moved into $B^{(0)}$ to make $\varepsilon_t \sim \mathcal{N}(\mathbf{0}, I)$, yielding

$$\begin{aligned} - \sum_{t=1}^T \log p(\varepsilon_t) &= \frac{nT}{2} \log(2\pi) + \frac{1}{2} \sum_{t=1}^T \|\varepsilon_t\|_2^2 \\ &= \frac{nT}{2} \log(2\pi) + \frac{1}{2} \|\mathcal{E}\|_F^2 \end{aligned} \quad (9)$$

for $\mathcal{E} = \varepsilon_{1:T}$. The constant is ignored in the optimization.

Third, rather than merely consider a maximum likelihood objective, we can consider the maximum a posteriori (MAP) estimate given by the introduction of a prior $\log p(\Theta)$ over the model parameters $\Theta = (A, B, \Sigma)$. Since the parameters A and B do not typically have distributional assumptions, we view the choice of priors rather as regularizers:

$$-\log p(\Theta) = -\log p(A) - \log p(B) = R(A) + G(B),$$

for convex functions R and G . Any convex regularizer on A is acceptable. The choice of regularizer on B is more subtle, since for any s , $B\mathcal{E} = (Bs^{-1})(s\mathcal{E})$: $G(B)$ is required to

prevent B from being scaled up, pushing $\|\mathcal{E}\|_F^2$ to zero. We consider $G(B) = \|B\|_F^2$ for $B = [B^{(0)}; \dots; B^{(q)}]$, which effectively controls the size of B and, importantly, also results in a global reformulation given in Theorem 2.

Finally, as noted, we can consider any natural exponential family distribution for $\boldsymbol{\eta}_t$ rather than merely assuming Gaussian. The negative log-likelihood for such a distribution corresponds to a regular Bregman divergence (see Appendix ??), allowing one to write the final estimation criterion in terms of a convex loss function $L(\cdot | \mathbf{x}_t)$ as

$$\begin{aligned} \min_{A, B, \mathcal{E}} \sum_{t=1}^T L \left(\sum_{i=1}^p A^{(i)} \mathbf{x}_{t-i} + \sum_{j=0}^q B^{(j)} \varepsilon_{t-j} \mid \mathbf{x}_t \right) \\ + \alpha \left(\|\mathcal{E}\|_F^2 + G(B) \right) + \gamma R(A), \end{aligned} \quad (10)$$

for regularization parameters α and γ .

Efficient parameter estimation

Although $L(\cdot | \mathbf{x}_t)$ is convex, (10) is not jointly convex due to the coupling between B and \mathcal{E} . However, using recent insights from matrix factorization (Bach, Mairal, and Ponce 2008; Dudik, Harchaoui, and Mallick 2012; White et al. 2012) one can reformulate (10) as a convex optimization.

For fixed autoregressive parameters, A , let $L_{A,t}(z) = L(z + \sum_{i=1}^p A^{(i)} \mathbf{x}_{t-i} \mid \mathbf{x}_t) + \gamma R(A)$, which is still convex in z .² By introducing the change of variables, $Z = B\mathcal{E}$, the optimization over B and \mathcal{E} given A can be written as

$$\begin{aligned} \min_{A, B, \mathcal{E}} \sum_{t=1}^T L_{A,t} \left(\sum_{j=0}^q B^{(j)} \varepsilon_{:,t-j} \right) + \alpha \left(\|\mathcal{E}\|_F^2 + G(B) \right) \\ = \min_{A, Z} \sum_{t=1}^T L_{A,t} \left(\sum_{j=0}^q Z_{:,t-j}^{(j)} \right) + \alpha \min_{\substack{B, \mathcal{E} \\ B\mathcal{E}=Z}} \left(\|\mathcal{E}\|_F^2 + G(B) \right). \end{aligned} \quad (11)$$

Surprisingly, this objective can be re-expressed in a convex form since

$$\|Z\| = \min_{B, \mathcal{E}: B\mathcal{E}=Z} \left(\|\mathcal{E}\|_F^2 + G(B) \right) \quad (12)$$

defines an induced norm on Z (established in Theorems 2 and 3 below) allowing (11) to be equivalently expressed as:³

$$\sum_{t=1}^T L \left(\sum_{i=1}^p A^{(i)} \mathbf{x}_{t-i} + \sum_{j=0}^q Z_{:,t-j}^{(j)} \mid \mathbf{x}_t \right) + \alpha \|Z\| + \gamma R(A)$$

Therefore, one can alternate between A and Z to obtain a globally optimal solution, then recover B and \mathcal{E} from Z , as discussed in the next section. Proofs for the following two theorems are provided in Appendix ??.

Theorem 2. *The regularized ARMA(p, q) estimation problem for $G(B) = \|B\|_F^2$ is equivalent to*

$$(11) = \min_{A, Z} \sum_{t=1}^T L_{A,t} \left(\sum_{j=0}^q Z_{:,t-j}^{(j)} \right) + 2\alpha \|Z\|_{tr}$$

with a singular value decomposition recovery: $Z = U\Sigma V^T$ giving $B = U\sqrt{\Sigma}$ and $\mathcal{E} = \sqrt{\Sigma}V^T$.

²Proved in Lemma ??, Appendix ?? for completeness

³Proved in Corollary ??, Appendix ?? for completeness.

The estimation problem is more difficult with the second choice of regularizer; to get an exact formulation, we need to restrict $q = 1$, giving $Z = \begin{bmatrix} B^{(0)} \\ B^{(1)} \end{bmatrix} \mathcal{E}$.

Theorem 3. *The regularized ARMA($p, 1$) estimation problem for $G(B) = \max_{j=0, \dots, q} \|B^{(j)}\|_F^2$ is equivalent to*

$$(11) = \min_{A, Z} \sum_{t=1}^T L_{A,t} \left(\sum_{j=0}^q Z_{:,t-j}^{(j)} \right) + \max_{0 \leq \rho \leq 1} \|W_\rho^{-1} Z\|_{tr} \quad (13)$$

where $W_\rho := \begin{bmatrix} 1/\sqrt{\rho} I_n & 0 \\ 0 & 1/\sqrt{1-\rho} I_n \end{bmatrix}$. Moreover $\|W_\rho^{-1} Z\|_{tr}$ is concave in ρ over $[0, 1]$, enabling an efficient line-search.

Identifiability and optimal parameter recovery

One desirable ideal for estimation is identifiability: being able to “identify” parameters uniquely. For a strictly convex loss function, L , the convex regularized ARMA optimization in (13) produces a unique moving average variable, Z . This identifiable matrix is sufficient for correctly estimating the autoregressive parameters for the ARMA model, which can be all that is required for forecasting in expectation.

It might be desirable, however, to recover the factors B and \mathcal{E} to gain further insight into the nature of the time series. Unfortunately, unlike Z , the factors that satisfy $B\mathcal{E} = Z$ are not unique. Worse, if one simply recovers any B and \mathcal{E} that satisfies $B\mathcal{E} = Z$, the recovered innovations \mathcal{E} need not be Gaussian distributed. This issue can be addressed via a careful recovery procedure that finds a particular pair B and \mathcal{E} with the same regularization penalty as Z . Let

$$\text{Factors}(Z) = \{(B, \mathcal{E}) : B\mathcal{E} = Z \text{ and } G(B) + \|\mathcal{E}\|_F^2 = \|Z\|\}$$

This set of solutions satisfies the desired distributional properties, but is invariant under scaling and orthogonal transformations: for any $(B, \mathcal{E}) \in \text{Factors}(Z)$, (i) for $s = G(B)/\|\mathcal{E}\|_F$, $(B(1/s), s\mathcal{E}) \in \text{Factors}(Z)$ and (ii) for any orthogonal matrix $P \in \mathbb{R}^{n \times n}$, $(BP, P'\mathcal{E}) \in \text{Factors}(Z)$ since the Frobenius norm is invariant under orthogonal transformations. When $G(B) = \|\mathcal{E}\|_F$, a solution from $\text{Factors}(Z)$ can be computed from the singular value decomposition of Z , as shown in Theorem 2. When $G(B) = \max_j \|B^{(j)}\|_F^2$, a boosting approach can be used for recovery; see Appendix ?? for details as well as a discussion on recovering Laplacian instead of Gaussian innovations.

Computational complexity

The overall estimation procedure is outlined in Algorithm 1 for $G(B) = \|B\|_F^2$; the approach is similar for the other regularizer, but with an outer line search over ρ . The computational complexity is governed by the matrix multiplication to compute the autoregressive and moving average components, and by the use of the singular value decomposition. The matrix multiplication for AX_p is $O(Tpn^2)$, which dominates the cost of computing the autoregressive loss, corresponding to ARloss in Algorithm 1. For the moving average loss, MALoss in Algorithm 1, the thin SVD of $Z \in \mathbb{R}^{qn \times T}$ has complexity $O(Tq^2n^2)$ and the multiplication of UV' is also $O(Tq^2n^2)$. Thus, each call to ARloss is $O(Tpn^2)$ and

Algorithm 1 RARMA(p, q)

Input: X, p, q, α, γ

Output: A, B, \mathcal{E}

```

1:  $X_p = \mathbf{0}$  // History matrix,  $X_p \in \mathbb{R}^{np \times T}$ 
2: for  $i = 1, \dots, p$ ,  $X_p(:, i) = [X(:, i-1); \dots; X(:, i-p)]$ 
3:  $[f, g] = \text{MALoss}(Z, A)$ 
4:  $[U, \Sigma, V] = \text{svd}(Z)$ 
5:  $Y = AX_p + \sum_{j=1}^q Z(j : n(j+1) - 1, :)$ 
6:  $f = L(Y; X) + \alpha \text{sum}(\text{diag}(\Sigma))$ 
7:  $g = \text{repmat}(\nabla_Y L(Y; X), q, 1)$ 
8: // Zero out unused parts of  $Z$ 
9: for  $j = 1, \dots, q$ ,  $g(j : n(j+1), (t-j+1) : t) = 0$ 
10:  $g = g + \alpha UV'$ 
11:  $[f, g] = \text{ARloss}(A, Z)$ 
12:  $Y = AX_p + \sum_{j=1}^q Z(j : n(j+1) - 1, :)$ 
13:  $f = L(Y; X) + \alpha R(A)$ 
14:  $g = (\nabla_Y L(Y; X))X_p + \gamma \nabla R(A)$ 
15: Initialize  $Z = \mathbf{0}, A = \mathbf{0}$ 
16: // Apply your favourite optimizer to the AR search
17:  $A = \text{lbfgs}(\text{ARloss}(\cdot, Z), A)$ 
18: // Iterate between  $A$  and  $Z$ 
19:  $[A, Z] = \text{iterate}(\text{ARloss}, \text{MALoss}, A, Z)$ 
20: // Recover the optimal  $B$  and  $\mathcal{E}$ 
21:  $[U, \Sigma, V] = \text{svd}(Z)$ 
Return:  $A, B = U\Sigma^{1/2}, \mathcal{E} = \Sigma^{1/2}V'$ 

```

each call to MALoss is $O(Tq^2n^2)$. The initial solution of A , which involves solving a basic vector autoregressive model, is $i_1 O(Tpn^2)$ where i_1 is the number of iterations (typically small). For i_2 the number of iterations between A and B : $\text{RARMA cost} = \text{VAR cost} + i_2(O(Tpn^2) + O(Tq^2n^2)) = (i_1 + i_2)O(Tpn^2) + i_2O(Tq^2n^2) \approx O(T(p + q^2)n^2)$.

Experimental evaluation

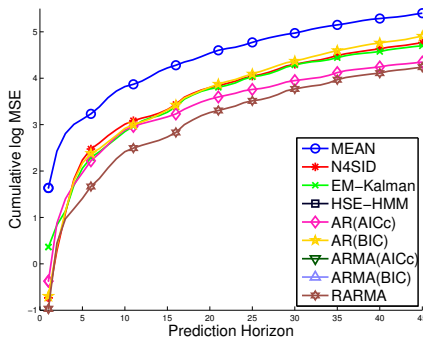
In this section, regularized ARMA is compared to a wide range of time series methods for the task of forecasting future observations on both synthetic and real-world data. As discussed in Appendix ??, forecasts are performed using only the autoregressive parameters. In the final section, there is also a comparison on estimating the underlying autoregressive parameters of RARMA using $q = 0$ versus $q > 0$.

Several algorithms are compared: MEAN, which uses the mean of the training sequence as the prediction, a popular subspace identification method (N4SID) (Van Overschee and De Moor 1994), expectation-maximization to learn the parameters for a Kalman filter (EM-Kalman), Hilbert space embeddings of hidden Markov models (HSE-HMM) (Song et al. 2010; Boots and Gordon 2012),⁴ maximum likelihood estimation of vector AR (AR), the Hannan-Rissanen method for ARMA (ARMA) (Hannan and Kavalieris 1984) and global estimation of regularized ARMA (RARMA). We also compared to local alternation of the RARMA objective

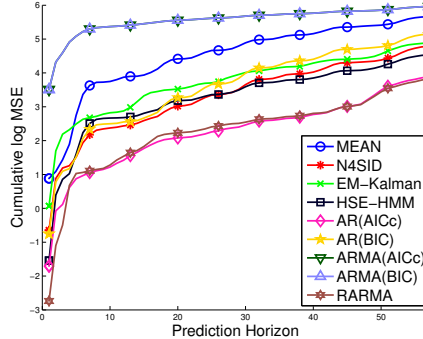
⁴In addition to the two method-of-moments approaches, N4SID and HSE-HMM, we tried a third state-space technique (Anandkumar, Hsu, and Kakade 2012), with no previous published empirical demonstrations. It performed poorly and so is omitted.

Table 1: For each dataset, the first column contains the test MSE (with standard error in parentheses) and the second the percentage of trials that were stable. The stability rates are measured using a threshold: eigenvalues $< 1 + \epsilon = 1.01$. The method(s) with the most t -test wins with significance level of 5% are bold for each dataset. Stable rates are key for iterated prediction performance; large MSE is mainly due to unstable trials.

ALGORITHM	N6-P2-Q2		N9-P3-Q3		N12-P3-Q3		N15-P3-Q3		N12-P4-Q4		CAC		ATLANTIC	
MEAN	2.85(0.13)	1.00	6.64(0.27)	1.00	12.5(0.44)	1.00	21.2(1.06)	1.00	6.81(0.19)	1.00	4.92(0.17)	1.00	5.03(0.39)	1.00
N4SID	3.23(0.21)	1.00	6.82(0.30)	1.00	12.9(0.58)	1.00	24.7(1.46)	1.00	6.85(0.19)	1.00	2.60(0.35)	1.00	2.10(0.98)	1.00
EM-KALMAN	4.27(0.30)	0.97	11.7(1.17)	0.94	19.0(1.19)	0.95	32.8(3.03)	0.89	15.9(1.45)	0.88	2.43(0.26)	1.00	2.33(0.66)	1.00
HSE-HMM	13.5(12.8)	0.95	1070(1469)	0.95	353(514)	0.95	2017(3290)	0.95	31.8(28.6)	0.91	7637(1433)	0.88	1.63(0.53)	1.00
AR(AICc)	1.83(0.29)	0.96	8.99(2.38)	0.88	16.9(2.69)	0.84	80.2(24.7)	0.69	24.8(29.6)	0.69	1.71 (0.31)	1.00	0.85 (0.37)	1.00
AR(BIC)	1.67(0.25)	0.97	6.42(1.27)	0.91	10.7(1.22)	0.93	34.0(5.20)	0.85	5.25(0.54)	0.82	3.00(0.27)	1.00	2.99(0.81)	1.00
ARMA(AICc)	1.63(0.2)	0.98	4.93(0.62)	0.93	8.52(0.73)	0.96	27.7(3.64)	0.86	5.49(0.48)	0.88	101(42.9)	1.00	6.80(2.85)	1.00
ARMA(BIC)	1.68(0.25)	0.98	4.81(0.58)	0.95	8.40(0.59)	0.97	29.4(5.76)	0.91	5.26(0.48)	0.97	107(49.3)	1.00	6.80(2.85)	1.00
RARMA	1.29 (0.08)	1.00	4.10 (0.30)	0.97	7.49 (0.41)	0.99	15.7 (0.92)	0.98	4.69 (0.21)	0.99	1.53 (0.27)	1.00	0.80 (0.35)	1.00



(a) CAC.



(b) Atlantic.

Figure 2: Cumulative test MSE in log scale on two real-world datasets. Each model is iterated for (a) 40 and (b) 60 steps, respectively. AR(BIC) and AR(AICc) have significantly different performance, indicating the importance in selecting a good lag length for AR. HSE-HMM is unstable for CAC, but performs reasonably well for Atlantic. The best performing methods are AR(AICc) and RARMA. RARMA has strong first step prediction in both cases and has very good early predictions in CAC.

(11); for both the best and worst random initializations, however, the results were always worse than global RARMA, slower by more than an order of magnitude and often produced unstable results. Therefore, these local alternator results are omitted, with the focus instead on the comparison between the RARMA objective and the other approaches. The HR method is used for the ARMA implementation, because a recent study (Kascha 2012) shows that HR is reasonably stable compared to other vector ARMA learning algorithms. A third step was added for further stability, in which the $A^{(i)}$ are re-learned from the observations with the MA component (from the second step) removed. The built-in Matlab implementations were used for AR and N4SID.

The lag parameters p and q were selected according to standard criteria in time series. For AR and ARMA, the parameters p and q are chosen using BIC and AICc, and reported separately due to interesting differences in their performance. For N4SID, the built-in Matlab implementation chose the best order. For RARMA, because of the temporal structure in the data, parameters were chosen by performing estimation on the first 90% of the training sequence and evaluating model performance on the last 10% of the training sequence. We use a robust loss, the Huber loss, for RARMA, which is easily incorporated due to the generality of RARMA. Autoregressive models can also easily use the Huber loss; we therefore directly compare RARMA to only using an autoregressive component in the last section.

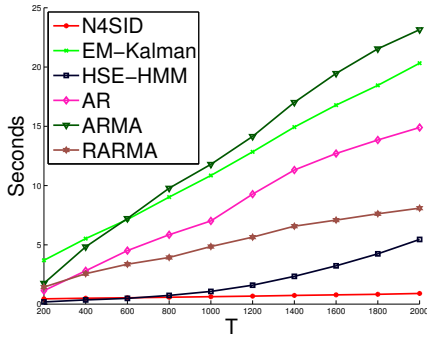
Synthetic experiments: Synthetic data sets are generated from an ARMA(p, q) model. An ARMA model is called unstable if the spectra of the AR matrices exceeds the unit circle on the complex plane (Lütkepohl 2007); intuitively, iterating a dynamics matrix $A = U\Sigma V'$ that has any $\Sigma(i, i) > 1$ for t steps, $A^t = U\Sigma^t V'$, is expansive. See Appendix ?? for details about the procedure for generating stable ARMA models. For each (p, q, n) configuration, 500 data sequences are generated, each with 300 samples partitioned into 200 training points followed by 100 test points.

Table 1 shows the results, including the test mean squared error (MSE) and the stability rates over all trials. RARMA with global optimization is the best model on each data set in terms of MSE. Learning is generally more difficult as the dimension increases, but RARMA performs well even when most algorithms fail to beat the baseline (MEAN) and maintains a reasonably stable rate.


Figure 3(a) illustrates a runtime comparison, in CPU seconds. The synthetic model is fixed to $n = 9$ and $p = q = 3$, with an increasing number of training points. RARMA is comparable to other methods in terms of efficiency and scales well with increasing number of samples T .

Experiments on real time series: To see how our method performs on real-world data, we experimented on two real-world datasets from IRI.⁵ These climate datasets consist of monthly sea surface temperature on the tropical Pacific


⁵<http://iridl.ldeo.columbia.edu/SOURCES/>



(a) Runtimes for increasing T



(b) Parameter recovery for $q = 0$ vs. $q > 0$



(c) Forecasting for $q = 0$ vs. $q > 0$

Figure 3: (a) Runtimes for an increasing number of moving samples, for multivariate series with $n = 9$ and $p = q = 3$. For (b) and (c), the relative error is reported between $\text{RARMA}(p, q)$ and the best $\text{RARMA}(m, 0)$ model, with $m = p, p + q$ (with results for each m separately in Appendix ??, Figure ??). The plot is symmetric, where at 4, $\text{RARMA}(p, q)$ has 4x lower error (good), and at 0.25, has 4x higher error (bad). The x-axis shows increasing lag and the y-axis increasing moving average variance. As the variance is increased beyond $\exp(-3)$, using the MEAN as a predictor begins to outperform all three methods and the moving average component begins to dominate the autoregressive component. For (b), the comparison is the ℓ_1 error between the recovered A parameters and the true parameters, cut-off at p for $\text{RARMA}(p + q, 0)$. As p, q increase, $\text{RARMA}(p + q, 0)$ becomes the dominant method for obtaining the underlying parameters, despite the fact that only the first p components of the learned A are used. For (c), the comparison is with forecasting accuracy for a horizon of 10, measured with ℓ_1 error. Using $q > 0$ clearly improves performance for lower variance levels; the mean predictor begins to outperform all three at e^{-2} .

Ocean (CAC) and the tropical Atlantic Ocean (Atlantic). The area size of CAC is 84×30 with 399 points, while the area of Atlantic is 22×11 with 564 points. We use the first 30×30 locations for CAC and the first 9×9 locations for Atlantic. These areas are further partitioned into grids of size 3×3 and vectorized to obtain observations $\mathbf{x}_t \in \mathbb{R}^9$. The data for each location is normalized to have sample mean of zero and sample standard deviation of one in the experiments. The first 90% of the sequence is used as training set, the last 10% as the test set.

Table 1 shows the test MSE and Figure 2 shows the cumulative MSE in log scale. As in the synthetic experiments, RARMA is consistently among the best models in terms of MSE. Moreover, when examining iterated forecasting accuracy in Figure 2, RARMA is better for early predictions (about the first 30 predictions) on the real datasets.

Investigating the moving average component: The final comparison is an investigation into the importance of the moving average component, versus simply using an AR model. $\text{RARMA}(p, q)$ is compared to $\text{RARMA}(p, 0)$ and $\text{RARMA}(p + q, 0)$ for two settings: recovering the true autoregressive parameters, A , and accuracy in forecasting. The same code is run for all three methods, simply with different p, q settings. The comparison is over increasing lag and increasing variance of the moving average component. The results are presented in Figure 3, with a more complete figure in Figure ??, Appendix ??. The heat map presents the relative error between $\text{RARMA}(p, q)$ and the best of $\text{RARMA}(p, 0)$ and $\text{RARMA}(p + q, 0)$; values greater than 1 indicate superior performance for $\text{RARMA}(p, q)$.

These results indicate two interesting phenomena. First, including the moving average component significantly improves forecasting performance when the variance is relatively small. As the variance reached levels where the MEAN began to outperform all three techniques, the models with $q = 0$ were slightly better. Second, RARMA with

$q > 0$ performs noticeably better for forecasting but typically performed about the same or worse for extracting the underlying autoregressive parameters. This result suggests that, if the ultimate goal is forecasting, we need not focus so much on identification. Importantly, because vector ARMA models can now be solved globally and efficiently, there is little disadvantage in using this more powerful model, and strong benefits in some cases.

Conclusion

This paper tackles a long-standing problem in time series modeling: tractable maximum likelihood estimation of multivariate ARMA models. The approach involves three key components: (1) estimating stochastic ARMA models, which relaxes the requirement that the innovations exactly equal the residuals, (2) characterizing the independent Gaussian structure of the innovations using a regularizer and (3) developing a theoretically sound convex reformulation of the resulting stochastic multivariate ARMA objective. Solving this convex optimization is efficient, guarantees global solutions and outperformed previous ARMA and state-space methods in forecasting on synthetic and real datasets.

These results suggest stochastic regularized ARMA is a promising direction for time series modeling, over conventional (deterministic) ARMA. Stochastic ARMA is similarly motivated by the Wold representation, but is amenable to optimization, unlike deterministic ARMA. Moreover, the regularized ARMA objective facilitates development of estimation algorithms under generalized innovations. Though the focus in this work was on a convex formulation for Gaussian innovations, it extends to Laplacian innovations for a $(2, 1)$ -block norm regularizer. Advances in optimizing structured norm objectives advance global estimation of regularized ARMA models for novel innovation properties.

Acknowledgements

This work was supported by NSERC, Alberta Innovates Technology Futures and the University of Alberta.

References

- Anandkumar, A.; Hsu, D.; and Kakade, S. M. 2012. A method of moments for mixture models and hidden Markov models. *arXiv:12030683v3 [cs.LG]*.
- Anava, O.; Hazan, E.; Mannor, S.; and Shamir, O. 2013. Online learning for time series prediction. In *Proceedings of the 26th Annual Conference on Learning Theory*, 1–13.
- Andersson, S. 2009. Subspace estimation and prediction methods for hidden Markov models. *The Annals of Statistics* 4131–4152.
- Arslan, O. 2010. An alternative multivariate skew Laplace distribution: properties and estimation. *Statistical Papers* 865–887.
- Bach, F.; Mairal, J.; and Ponce, J. 2008. Convex sparse matrix factorizations. *arXiv:0812.1869v1 [cs.LG]*.
- Banerjee, A.; Merugu, S.; Dhillon, I. S.; and Ghosh, J. 2005. Clustering with Bregman divergences. *Journal of Machine Learning Research* 6:1705–1749.
- Bauer, D. 2005. Estimating linear dynamical systems using subspace methods. *Econometric Theory* 21(01).
- Benjamin, M. A.; Rigby, R. A.; and Stasinopoulos, D. M. 2003. Generalized autoregressive moving average models. *Journal of the American Statistical Association* 98(461):214–223.
- Benveniste, A.; Metivier, M.; and Priouret, P. 2012. *Adaptive Algorithms and Stochastic Approximations*. Springer.
- Boots, B., and Gordon, G. 2012. Two-manifold problems with applications to nonlinear system identification. In *Proceedings of the 29th International Conference on Machine Learning*, 623–630.
- Brockwell, P. J., and Davis, R. A. 2002. *Introduction to Time Series and Forecasting*. Springer.
- Brown, P. F.; Pietra, V. J. D.; Pietra, S. A. D.; and Mercer, R. L. 1993. The mathematics of statistical machine translation: parameter estimation. *Computational Linguistics* 19(2):263–311.
- Candes, E. J.; Li, X.; Ma, Y.; and Wright, J. 2011. Robust principal component analysis? *Journal of the ACM* 58:1–37.
- Cramér, H. 1946. *Mathematical Methods of Statistics*. Princeton University Press.
- Dudik, M.; Harchaoui, Z.; and Malick, J. 2012. Lifted coordinate descent for learning with trace-norm regularization. *Proceedings of the 15th International Conference on Artificial Intelligence and Statistics* 22:327–336.
- Durbin, J. 1960. The fitting of time-series models. *Revue de l'Institut International de Statistique* 28:233–244.
- Foster, D. P.; Rodu, J.; and Ungar, L. H. 2012. Spectral dimensionality reduction for HMMs. *arXiv:1203.6130v1*.
- Hannan, E. J., and Kavalieris, L. 1984. Multivariate linear time series models. *Advances in Applied Probability* 16:492–561.
- Hsu, D.; Kakade, S.; and Zhang, T. 2012. A spectral algorithm for learning Hidden Markov Models. *Journal of Computer and System Sciences* 1460–1480.
- Kascha, C. 2012. A comparison of estimation methods for vector Autoregressive Moving-Average Models. *Econometric Reviews* 31(3):297–324.
- Katayama, T. 2006. *Subspace Methods for System Identification*. Springer.
- Lütkepohl, H. 2007. *New Introduction to Multiple Time Series Analysis*. Springer.
- Mauricio, J. A. 1995. Exact maximum likelihood estimation of stationary vector ARMA models. *Journal of the American Statistical Association* 282–291.
- Moonen, M., and Ramos, J. 1993. A subspace algorithm for balanced state space system identification. *IEEE Transactions on Automatic Control* 38(11):1727–1729.
- Neal, R. M., and Hinton, G. 1998. A view of the EM algorithm that justifies incremental, sparse, and other variants. *Learning in Graphical Models* 355–368.
- Scargle, J. D. 1981. Studies in astronomical time series analysis. *The Astrophysical Journal Supplement Series* 835–853.
- Shah, P.; Bhaskar, B. N.; Tang, G.; and Recht, B. 2012. Linear system identification via atomic norm regularization. *arXiv:1204.0590v1 [math.OC]*.
- Song, L.; Boots, B.; Siddiqi, S.; Gordon, G.; and Smola, A. 2010. Hilbert space embeddings of hidden Markov models. In *Proceedings of the 27th International Conference on Machine Learning*, 991–998.
- Thiesson, B.; Chickering, D. M.; Heckerman, D.; and Meek, C. 2012. ARMA time-series modeling with graphical models. In *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence*, 552–560.
- Van Overschee, P., and De Moor, B. 1994. N4SID: Subspace algorithms for the identification of combined deterministic-stochastic systems. *Automatica* 30(1):75–93.
- Viberg, M. 1995. Subspace-based methods for the identification of linear time-invariant systems. *Automatica* 31(12):1835–1851.
- White, M.; Yu, Y.; Zhang, X.; and Schuurmans, D. 2012. Convex multi-view subspace learning. In *Advances in Neural Information Processing Systems*, 1673–1681.
- Wiesel, A.; Bibi, O.; and Globerson, A. 2013. Time varying autoregressive moving average models for covariance estimation. *IEEE Transactions on Signal Processing* 61(11):2791–2801.
- Wold, H. O. A. 1938. *A Study in The Analysis of Stationary Time Series*. Almqvist & Wiskell.
- Zhang, X.; Yu, Y.; White, M.; Huang, R.; and Schuurmans, D. 2011. Convex sparse coding, subspace learning, and semi-supervised extensions. In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*, 567–573.
- Zhao, H., and Poupart, P. 2014. A sober look at spectral learning. *arXiv:1406.4631v1 [cs.LG]*.