

Spatial Projection of Multiple Climate Variables Using Hierarchical Multitask Learning

André R. Gonçalves

Center for Research and Development
in Telecommunication (CPqD), Brazil
andreg@cpqd.com.br

Arindam Banerjee

Dept. of Comp. Sci. & Engg.
University of Minnesota, USA
banerjee@cs.umn.edu

Fernando J. Von Zuben

School of Elec. and Comp. Eng.
University of Campinas, Brazil
vonzuben@dca.fee.unicamp.br

Abstract

Future projection of climate is typically obtained by combining outputs from multiple Earth System Models (ESMs) for several climate variables such as temperature and precipitation. While IPCC has traditionally used a simple model output average, recent work has illustrated potential advantages of using a multitask learning (MTL) framework for projections of individual climate variables. In this paper we introduce a framework for hierarchical multitask learning (HMTL) with two levels of tasks such that each *super-task*, i.e., task at the top level, is itself a multitask learning problem over *sub-tasks*. For climate projections, each super-task focuses on projections of specific climate variables spatially using an MTL formulation. For the proposed HMTL approach, a group lasso regularization is added to couple parameters across the super-tasks, which in the climate context helps exploit relationships among the behavior of different climate variables at a given spatial location. We show that some recent works on MTL based on learning task dependency structures can be viewed as special cases of HMTL. Experiments on synthetic and real climate data show that HMTL produces better results than decoupled MTL methods applied separately on the super-tasks and HMTL significantly outperforms baselines for climate projection.

Introduction

Future projections of climate variables such as temperature, precipitation, and pressure are usually produced by physics-based models known as Earth System Models (ESMs). ESMs consist of four components and their interactions, viz. atmosphere, oceans, land, and sea ice (Tebaldi and Knutti 2007; IPCC 2013). Climate projections generated from ESMs form the basis for understanding and inferring future climate change, global warming, greenhouse gas concentration and its impact on Earth systems and other complex phenomena such as El Niño Southern Oscillation (ENSO). ENSO, for instance, has a global impact, ranging from droughts in Australia and northeast Brazil to heavy rains over Malaysia, the Philippines, and Indonesia (IPCC 2013). Then, producing accurate projections of climate variables is essential to anticipate extreme events.

Many ESMs have been developed by climate research institutes. A single and possibly more robust projection can be

built as a combination (ensemble) of multiple ESMs simulations (Tebaldi and Knutti 2007; McQuade and Monteoloni 2013; Sanderson, Knutti, and Caldwell 2015).

Recently, the problem of constructing ESMs ensemble was approached from a multitask learning (MTL) perspective (Gonçalves et al. 2014), where building an ESMs ensemble for each geographical location was viewed as a learning task. The joint estimation of the ESM ensemble produced more accurate projections than when independent estimation was performed for each location. The MTL method was able to capture the relationship among geographical locations (tasks) and used it to guide information sharing among tasks.

Modeling task relationship in multitask learning has been the focus of recent research (Zhang and Schneider 2010; Zhang and Yeung 2010; Yang, Li, and Zhang 2013; Gonçalves, Von Zuben, and Banerjee 2016). This is a fundamental step to promote information sharing only among related tasks, while avoiding the unrelated ones. Besides estimating task specific parameters (Θ), the task dependency structure (Ω) is also estimated from the data. The latter is usually estimated from the former, i.e., task dependency is defined based on the relation of the task parameters. Two tasks are said to be related if their parameters are related in some sense.

Inconsistently estimated task dependency structure in MTL can misguide information sharing and, hence, can be harmful to the MTL method performance. The problem of estimating statistical dependency structure of a set of random variables is known as structure learning (Meinshausen and Bühlmann 2006). Existing methods for the problem guarantee to recover the true underlying dependence structure given a sufficient amount of data samples. In the MTL case (Gonçalves, Von Zuben, and Banerjee 2016), the random variables are tasks parameters and, depending on the ratio between dimensionality and the number of tasks, the amount of data samples may not be sufficient.

In this paper, we introduce the *Hierarchical Multitask Learning* (HMTL) method that jointly learn multiple tasks by letting each task be, by itself, a MTL problem. The associated hierarchical multitask learning problem is then said to be composed of *super-tasks* and the tasks involved in each super-task as *sub-tasks*. Our formulation is motivated by the problem of constructing ESM ensembles for multiple cli-

mate variables, with multiple geolocations for each variable. The problem of obtaining ESMs weights for all regions for a certain climate variable is a super-task.

The paper is organized as follows. We first discuss existing MTL methods with task dependency estimation, on which our proposed ideas are built; then, we briefly talk about the climate projection problem that motivated our work. In the sequel, we present our proposed framework, perform experiments and analyze the results. Concluding remarks and next research steps complete the paper.

Notation. Let T be the number of super-tasks, d the problem dimension, and $n_{(t,k)}$ the number of samples for the (t,k) -th sub-task. For the purposes of the current paper, we assume that all super-tasks have m sub-tasks and all sub-tasks have problem dimension d . $X^{(t,k)} \in \mathbb{R}^{n_{(t,k)} \times d}$ and $\mathbf{y}^{(t,k)} \in \mathbb{R}^{n_{(t,k)} \times 1}$ are the input and output data for the k -th sub-task of the t -th super-task. $\Theta^{(t)} \in \mathbb{R}^{d \times m}$ is the matrix whose columns are the set of weights for all sub-tasks for the t -th super-task, that is, $\Theta^{(t)} = [\boldsymbol{\theta}^{(t,1)}, \dots, \boldsymbol{\theta}^{(t,m)}]$. For the ease of exposition, we represent $\{X\} = X^{(t,k)}$ and $\{Y\} = \mathbf{y}^{(t,k)}$, $k = 1, \dots, m_t$; $t = 1, \dots, T$. For the weight and precision matrices, $\{\Theta\} = \Theta^{(t)}$ and $\{\Omega\} = \Omega^{(t)}$, $\forall t = 1, \dots, T$. Identity matrix of dimension $m \times m$ is denoted by \mathbf{I}_m . $\mathcal{U}(a, b)$ is a uniform probability distribution in the range $[a, b]$.

Multitask Learning with Task Dependence Estimation

Explicitly modeling task dependencies has been made by means of Bayesian models. Features across tasks (rows of the parameter matrix Θ) were assumed to be drawn from a multivariate Gaussian distribution. Task relationship is then encoded in the inverse of the covariance matrix $\Sigma^{-1} = \Omega$, also known as *precision matrix*. Sparsity is desired in such matrix, as zero entries of the precision matrix indicate conditional independence between the corresponding two random variables (tasks) (Friedman, Hastie, and Tibshirani 2008). The associated learning problem (1) consists of jointly estimating the task parameters Θ and the precision matrix Ω , which is done by an alternating optimization procedure.

$$\begin{aligned} \min_{\Theta, \Omega} \quad & \sum_{k=1}^m L(X_k, \mathbf{y}_k, \Theta) - \log |\Omega| + \lambda_0 \text{tr}(\Theta \Omega \Theta^\top) + R(\Theta, \Omega) \\ \text{s.t.} \quad & \Omega \succeq 0. \end{aligned} \quad (1)$$

Note that in (1) the regularization penalty $R(\Theta, \Omega)$ is a general penalization function that will be discussed later in the paper. A solution for (1) alternates between the following two steps until a stopping criterion is met:

Step 1 Estimate Θ from current estimation of Ω ;

Step 2 Estimate Ω from updated parameters Θ .

Setting initial Ω to identity matrix, i.e., all tasks are independent at the beginning, is usually a suitable start.

In Step 1, task dependency information is incorporated into the joint cost function through the trace term

penalty $\text{tr}(\Theta \Omega \Theta^\top)$. It helps to promote information exchange among tasks. The problem associated with Step 2, known as *sparse inverse covariance selection problem* (Friedman, Hastie, and Tibshirani 2008), seeks to find some sparsity pattern in the precision matrix. Experimental analysis have shown that these approaches usually outperform MTL with pre-defined task dependency structure for a variety of problems (Zhang and Schneider 2010; Gonçalves et al. 2014).

Mathematical Formulation of ESMs Climate Projection

A common projection method is to perform the combination of multiple ESMs in a least square sense, that is, to estimate a set of weights for the ESMs based on past observations.

For a given location k the predicted climate variable (temperature, for example) for a certain timestamp i (expected mean temperature for a certain month/year, for example) is given by:

$$\hat{y}_k^i = \langle \mathbf{x}_k^i, \boldsymbol{\theta}_k \rangle + \epsilon_k^i \quad (2)$$

where \mathbf{x}_k^i is the set of values predicted by the ESMs for the k -th location in the timestamp i , $\boldsymbol{\theta}_k$ is the weight vector of each ESM for the k -th location, and ϵ_k^i is a residual. The weight vector $\boldsymbol{\theta}_k$ is estimated from a training data. The combined estimate \hat{y}_k^i is then used as a more robust prediction of temperature for the k -th location in a certain month/year in the future.

ESMs weights ($\boldsymbol{\theta}_k$) are defined for each geographical location and it possibly varies for different locations. It is possible that some ESMs are more accurate for some regions/climate than others and the difference between weights of two locations will reflect this behavior. In summary, the ESMs ensemble consists of solving a least square problem for each geographical location.

The ESMs weights may vary for the prediction of different climate variables, such as precipitation, temperature, and pressure. Then, solving an MTL problem for each climate variable is required. In this paper, we propose to simultaneously tackle multiple MTL problems through a hierarchical (two-level) MTL formulation.

The HMTL Formulation

The HMTL formulation seeks to minimize the following cost function $C(\Gamma)$ with $\Gamma = \{\{\Theta\}, \{\Omega\}, \lambda_0\}$:

$$\begin{aligned} \mathcal{L}(\Gamma) = \sum_{t=1}^T \left(\sum_{k=1}^{m_t} L(X^{(t,k)} \boldsymbol{\theta}^{(t,k)}, \mathbf{y}^{(t,k)}) \right. \\ \left. - \log |\Omega^{(t)}| + \lambda_0 \text{tr}(S^{(t)} \Omega^{(t)}) \right) + \mathcal{R}(\{\Omega\}) \end{aligned} \quad (3)$$

where $\mathcal{R}(\{\Omega\})$ is a regularization term over the precision matrices, $S^{(t)}$ is the sample covariance matrix of the task parameters for the t -th super-task. For simplicity, we dropped the ℓ_1 -penalization on the weight matrix Θ as is often done in MTL (Zhang and Schneider 2010; Yang, Li, and Zhang 2013; Gonçalves, Von Zuben, and Banerjee 2016). However, it can be added, if one desires sparse Θ , with minor changes

in the algorithm. All super-tasks are assumed to have the same number of sub-tasks. Squared loss was used as loss function.

The formulation (3) is a penalized cumulative cost function of the form (1) for several multitask learning problems. The penalty function $R(\{\Omega\})$ is to favor common structural sparseness across the precision matrices. Here, we focus on the group lasso penalty (Yuan and Lin 2006), which we denote by R_G , and is defined as

$$R_G(\{\Omega\}) = \lambda_1 \sum_{t=1}^T \sum_{k \neq j} |\Omega_{kj}^{(t)}| + \lambda_2 \sum_{k \neq j} \sqrt{\sum_{t=1}^T \Omega_{kj}^{(t)2}} \quad (4)$$

where λ_1 and λ_2 are two nonnegative tuning parameters. The first penalty term is an ℓ_1 -penalization of the off-diagonal elements, so that non-structured sparsity in the precision matrices is enforced. The larger the value of λ_1 , the sparser the precision matrices. The second term of the group sparsity penalty encourages the precision matrices for different super-tasks to have the same sparsity pattern. Group lasso does not impose any restriction on the degree of similarity of the non-zero entries of the precision matrices.

Setting λ_2 to zero, super-tasks are decoupled into independent MTL formulations. Then, λ_2 can be seen as a coupling parameter, as larger values push the super-tasks to be coupled, so that different $\Omega^{(t)}$ have similar sparsity patterns. On the other hand, lower (or zero) values lead to decoupled super-tasks, recovering existing MTL formulations, such as in (Gonçalves et al. 2014).

Figure 1 shows the hierarchy of tasks for the projection of multiple climate variables. At the level of super-tasks, group lasso regularization encourages precision matrices to have a similar sparseness pattern. The learned precision matrices are consequently used to control with whom each sub-task will share information.

The correspondence between the variables in the HMTL formulation and the elements in the climate problem is shown in Table 1.

Optimization

Optimization problem (3) is not jointly convex on $\{\Theta\}$ and $\{\Omega\}$, particularly due to the *trace* term which involves both variables. We then use an alternating minimization, in which $\{\Theta\}$ is held fix and optimize for $\{\Omega\}$ (we call it Ω -step), and similarly fix $\{\Omega\}$ and optimize for $\{\Theta\}$ (we call it Θ -step). Both steps now consist of convex problems, for which efficient methods have been proposed. In the experiments, 20 to 30 iterations were required for convergence.

Solving Θ -step The convex problem associated with this step is defined as

$$\min_{\{\Theta\}} \sum_{t=1}^T \sum_{k=1}^{m_k} L(X^{(t,k)} \theta^{(t,k)}, \mathbf{y}^{(t,k)}) + \lambda_0 \text{tr}(S^{(t)} \Omega^{(t)}). \quad (5)$$

Considering the squared loss function, Θ -step consists of two quadratic terms, as $\{\Omega\}$ are positive semidefinite matrices. Note that the optimization for each super-task weight

Table 1: Correspondence between HMTL variables and the components in the joint ESMs ensemble for the multiple climate variables problem.

Var.	HMTL meaning	Climate meaning
T	# of super-tasks	# of climate variables
m	# of sub-tasks in the t -th super-task	# of locations (equal for all climate variables)
$X^{(t,k)}$	data input for the k -th sub-task in the t -th super-task	ESMs predictions for the t -th climate variable in the k -th location
$\mathbf{y}^{(t,k)}$	data output for the k -th sub-task in the t -th super-task	observed values of the t -th climate variable in the k -th location
$\theta^{(t,k)}$	linear regression parameters for the k -th sub-task of the t -th super-task	ESMs weights for the t -th climate variable in the k -th location
$\Omega^{(t)}$	precision matrix for the t -th super-task	dependence among the ESMs weights for all locations for the t -th climate variable

matrix $\Theta^{(t)}$ are independent and can be performed in parallel. We used the L-BFGS (Liu and Nocedal 1989) method in the experiments.

Solving Ω -step The Ω -step is to solve the following optimization problem

$$\min_{\{\Omega\}} \sum_{t=1}^T \left(-\log |\Omega^{(t)}| + \lambda_0 \text{tr}(S^{(t)} \Omega^{(t)}) \right) + R_G(\{\Omega\}) \quad (6)$$

s.t. $\Omega^{(t)} \succeq 0, \forall t = 1, \dots, T.$

This step corresponds to the problem of joint learning multiple Gaussian graphical models and has been recently studied (Honorio and Samaras 2010; Danaher, Wang, and Witten 2014; Mohan et al. 2014). These formulations seek to minimize the penalized joint negative log likelihood in the form of (6) and they basically differ in the penalization term $R(\{\Omega\})$. Researchers have shown that the graphical models jointly estimated were able to increase the number of edges correctly identified while reducing the number of edges incorrectly identified, when compared to those independently estimated. An alternating direction method of multipliers (ADMM) proposed in (Danaher, Wang, and Witten 2014) was used to solve problem (6).

Algorithm (1) presents the pseudo-code for the proposed HMTL algorithm.

Related Works

A lattice graph is used to represent the regional relationship in (Subbian and Banerjee 2013), where immediate neighbor locations are assumed to have similar ESMs weights. Weights for all geolocations are estimated jointly (in a least square sense) with graph Laplacian regularization to encourage spatial smoothness.

In (McQuade and Monteleoni 2013) the online ESMs ensemble problem is tackled by using a lattice Markov Ran-

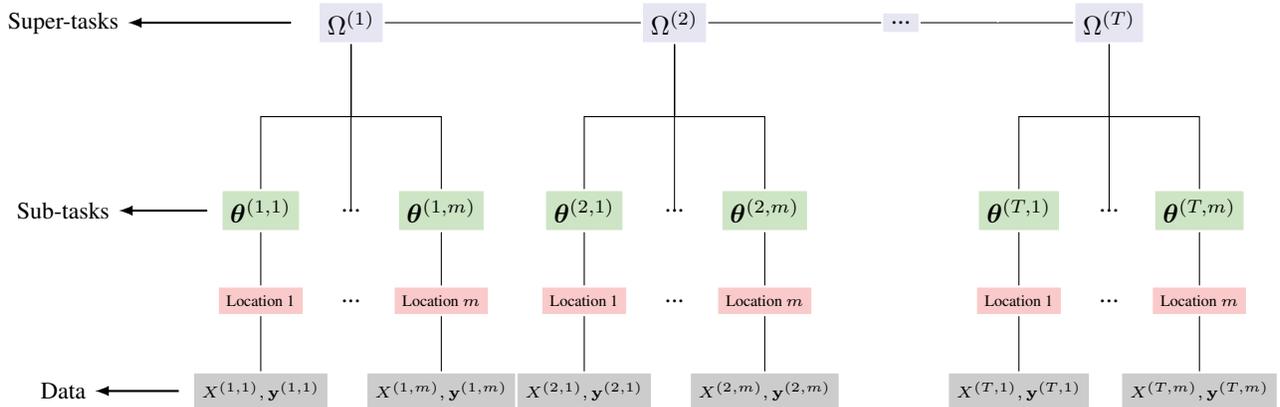


Figure 1: Hierarchy of tasks and their connection to the climate problem. Each super-task is a multitask learning problem for a certain climate variable, while sub-tasks are least square regressors for each geographical location.

Algorithm 1: HMTL algorithm.

Data: $\{\mathbf{X}\}, \{\mathbf{Y}\}$.
Input: $\lambda_0 > 0, \lambda_1 > 0$ and $\lambda_2 > 0$.
Result: $\{\Theta\}, \{\Omega\}$.

- 1 **begin**
- 2 $\Omega^{(t)} = \mathbf{I}_{m_t}, \forall t = 1, \dots, T$.
- 3 $\Theta^{(t)} = \mathcal{U}(-0.5, 0.5), \forall t = 1, \dots, T$.
- 4 **repeat**
- 5 Update $\{\Theta\}$ by solving (5);
- 6 Update $\{\Omega\}$ by solving (6);
- 7 **until** *stopping condition met*

dom Field. The state of each hidden variable, which is associated with a geographical location, is the identity of the best ESM for that specific location. The marginal probabilities of the hidden variables act as the weights of the ensemble. Hence, ESMs with a higher probability of being the best has a larger weight in the ensemble. At each time step marginal probabilities are updated based on the performance of the ESMs in the previous time step via the loopy belief propagation algorithm.

Differently, (Gonçalves et al. 2014) do not assume any fixed dependence graph, but, instead, estimate it through a multitask learning model that makes use of a sparse Gaussian Markov Random Field (GMRF) to capture the relationship of ESMs weights among locations. ESM weights for all locations and the GMRF are jointly estimated via an alternating minimization scheme.

Our work differs from the existing research as it (1) leverages information not only from immediate neighbors but also from any related geographical locations; (2) allows two levels of information sharing: ESMs weights and precision matrices that encodes the relationship of locations; and (3) handles the projection of multiple climate variables (exploring their resemblance) simultaneously.

Experiments

In this section we present experiments to compare the proposed HMTL with existing methods in the literature for both synthetic and real climate data.

Synthetic Data

We first generated a synthetic dataset to assess the performance of HMTL over traditional MTL methods. For comparison, we used the MTL method proposed in (Gonçalves et al. 2014; Gonçalves, Von Zuben, and Banerjee 2016), called MSSL, that has shown to be competitive with existing MTL algorithms including (Zhang and Yeung 2010; Kumar and Daume III 2012; Kang, Grauman, and Sha 2011). We also seek to investigate the effect of the increase in the number of super-tasks. For this analysis we generated 7 super-tasks containing 15 sub-tasks each, with dimension of 50. For each sub-task 100 data samples were generated. Inverse covariance matrices $\Omega^{(t)}, t = 1, \dots, T$ were drawn from a Wishart distribution with a scale matrix Λ and $n = 10$ degrees of freedom. Scale matrix Λ was designed to reflect a structure containing three groups of related variables. As sampled precision matrices are also likely to present the group structure, jointly sparse precision matrices are suitable to capture such pattern, which is precisely what HMTL formulation assumes.

Given $\Omega^{(t)}$, the sub-tasks parameters $\Theta^{(t)}$ were constructed as: $\Theta^{(t)}(j, :) = \mathcal{N}(0, \Omega^{(t)})$, $j = 1, \dots, d$; and, finally, the design matrix $X^{(t,k)}$ were sampled from $\mathcal{N}(0, I)$, and $Y^{(t,k)} = X^{(t,k)}\Theta^{(t)} + \epsilon^{(t,k)}$, where $\epsilon^{(t,k)} \sim \mathcal{N}(0, 0.1)$, $\forall k = 1, \dots, m; t = 1, \dots, T$.

Ten independent runs for both HMTL and MSSL were carried out. Each run with a different random train/test data split. Table 2 shows the relative improvement¹ of HMTL over MSSL for distinct number of super-tasks. Two scenarios were tested: using 50 and 30 samples for training and

¹Relative improvement is given by the difference between MSSL and HMTL performance (RMSE) divided by MSSL performance as percents (%).

the remaining for test (50 and 70, respectively). Penalization parameters (λ 's) were chosen by cross-validation on the training set.

Table 2: Average relative improvement (in %) of HMTL over MSSL for the synthetic dataset. HMTL produced lower RMSE as the number of super-tasks increase.

Training samples	# of super-tasks					
	2	3	4	5	6	7
50	2.2	4.0	4.9	5.5	5.2	5.8
30	22.3	23.4	25.1	25.7	24.2	24.3

From Table 2, we observe that HMTL has a continuous increase on the relative improvement compared to independently running MSSL, as the number of super-tasks grows. From the sixth super-task on, the relative improvement shows a tendency to stabilize. Such behavior is also observed in the scenario with only 30 training samples. It is worth noting that the smaller a set of training data (30 samples) the sharper the improvements obtained by HMTL over MSSL.

Climate Data

We collected monthly land temperature and precipitation data of 32 CMIP5 ESMs (Taylor, Stouffer, and Meehl 2012), from 1901 to 2000, in South America. Observed data provided by (Willmott and Matsuura 2001) was used. ESMs predictions and observed values from 250 locations in South America (distributed in a grid shape) were considered. From the HMTL perspective, the problem involves: two super-tasks, 250 sub-tasks (per super-task) with dimensionality of 32.

In the climate domain, it is common to work with the relative measure of the climate variable to a value of reference, which is obtained from past information. In our experiments, we directly work on the raw data (not detrended). We investigate the performance of the algorithm in both seasonal and annual time scales, with focus on winter and summer. All ESMs and observed data are in the same time and spatial resolution. Temperature is in degree Celsius and precipitation in cm.

Experimental Setup: Based on climate data from a certain past (training) period, model parameters are estimated and the inference produces its projections for the future (test). Clearly, the length of the training period affects the performance of the algorithm. A moving window of 20, 30 and 50 years were used for training and the next 10 years for test. The performance is measured in terms of root-mean-squared error (RMSE).

Seasonality strongly affect climate data analysis. Winter and summer precipitation patterns, for example, are distinct. Also, by looking at seasonal data, it becomes easier to identify anomalous patterns, possibly useful to characterize climate phenomena as El Niño. We extracted summer and winter data and performed climate variable projection specifically for these seasons.

Five baseline algorithms were considered:

1. **multi-model average** (MMA): set equal weights for all ESMs. This is currently performed by IPCC (IPCC 2013);
2. **best-ESM** in training phase: it is not an ensemble, but a single best ESM in terms of mean squared error;
3. **ordinary least square** (OLS): perform independent OLS for each location and climate variable;
4. **S²M²R** (Subbian and Banerjee 2013): can be seen as an MTL method with pre-defined location dependence given by the graph Laplacian. It incorporates spatial smoothing on ESMs weights.
5. **MSSL** (Gonçalves et al. 2014): run MSSL for each climate variable projection independently. The parameter-based version (*p*-MSSL) was used.

All the penalization parameters of the methods (λ 's in MSSL and HMTL) were chosen by cross-validation. From the training set, we selected the first 80% for training and the next 20% for validation. The best values in the validation set were selected. For example, in the scenario with 20 years of measurements for training, we took the first 16 years to re-ally train the model, and the next 4 years to analyze the performance of the method using a specific setting of λ 's. Using this protocol, the selected parameter values were: S²M²R used $\lambda = 1000$; MSSL $\lambda_0 = 0.1$ and $\lambda_1 = 0.1$; and HMTL $\lambda_0 = 0.1$, $\lambda_1 = 0.0002$, $\lambda_2 = 0.01$.

Results: Table 3 shows the RMSE of the projections produced by the algorithms and the observed values, for precipitation and temperature. First, we note that simply assigning equal weights to all ESMs does not seem to exploit the potential of ensemble methods. MMA (as used by IPCC, (IPCC 2013)) presented the largest RMSE among the algorithms for the majority of periods (*summer*, *winter* and *year*) and number of years for training. Second, the MTL methods, MSSL and HMTL, clearly outperform the baseline methods. S²M²R does not always produce better projections than OLS. In fact, it is slightly worse for the *year* dataset. As expected, the assumption of spatial neighborhood dependence does not seem to hold for all climate variables.

HMTL presented better results than performing MSSL for precipitation and temperature independently in many situations. HMTL was able to significantly reduce RMSE in *summer* precipitation projections, which has shown to be the most challenging scenario. Significant improvement was also seen for *winter* and *year* temperature projections.

RMSE per geographical location for precipitation and temperature are presented in Figure 2. For precipitation, more accurate projections (lower RMSE) was obtained by HMTL in Northernmost regions of South America, including Colombia and Venezuela. More accurate temperature projections were obtained in central North region of South America, which comprises part of the Amazon rainforest.

We believe that the improvements of the hierarchical MTL model are due to the same reasons as general MTL models: reduction of sample complexity by leveraging information from other related super-tasks. As a consequence, the precision matrices, which guide information sharing among sub-tasks, are estimated more accurately. The reduced sample complexity probably explains the better climate vari-

	# of training years	Best-ESM	OLS	S ² M ² R	MMA	MSSL	HMTL
Precipitation							
Summer	20	7.88 (0.44)	9.08 (0.54)	7.33 (0.68)	8.95 (0.27)	7.16 (0.43)	6.48 (0.34)**
	30	7.95 (0.55)	7.87 (0.63)	7.39 (0.86)	8.96 (0.26)	6.86 (0.48)	6.37 (0.29)*
	50	8.30 (0.71)	7.84 (1.13)	7.86 (1.12)	9.03 (0.30)	6.89 (0.55)	6.42 (0.33)
Winter	20	4.83 (0.26)	5.62 (0.30)	4.58 (0.39)	5.44 (0.24)	3.98 (0.21)	3.83 (0.22)
	30	4.86 (0.29)	4.83 (0.27)	4.68 (0.38)	5.41 (0.25)	3.94 (0.17)	3.80 (0.21)*
	50	4.92 (0.38)	4.64 (0.63)	4.77 (0.52)	5.33 (0.18)	3.84 (0.21)	3.70 (0.20)
Year	20	7.38 (0.17)	6.03 (0.65)	6.49 (0.49)	7.78 (0.14)	5.79 (0.16)	5.70 (0.16)
	30	7.41 (0.18)	6.21 (0.80)	6.57 (0.61)	7.76 (0.14)	5.72 (0.16)	5.66 (0.18)
	50	7.47 (0.26)	6.56 (1.07)	6.87 (0.80)	7.73 (0.14)	5.69 (0.23)	5.61 (0.22)
Temperature							
Summer	20	1.39 (0.23)	1.22 (0.10)	0.95 (0.13)	1.95 (0.02)	0.82 (0.08)	0.81 (0.01)
	30	1.47 (0.30)	1.21 (0.15)	1.09 (0.17)	1.96 (0.01)	0.84 (0.07)	0.80 (0.01)
	50	1.63 (0.35)	1.40 (0.19)	1.36 (0.20)	1.98 (0.01)	0.88 (0.05)	0.83 (0.01)*
Winter	20	1.58 (0.19)	1.48 (0.08)	1.18 (0.12)	2.08 (0.01)	1.03 (0.04)	1.02 (0.03)
	30	1.64 (0.26)	1.40 (0.13)	1.27 (0.16)	2.09 (0.01)	1.01 (0.04)	0.99 (0.03)
	50	1.77 (0.31)	1.55 (0.17)	1.51 (0.18)	2.08 (0.01)	1.04 (0.02)	0.98 (0.03)**
Year	20	1.64 (0.18)	1.10 (0.13)	1.13 (0.12)	2.11 (0.01)	1.00 (0.04)	0.91 (0.02)**
	30	1.70 (0.24)	1.20 (0.17)	1.24 (0.17)	2.12 (0.01)	1.00 (0.04)	0.91 (0.02)**
	50	1.83 (0.28)	1.47 (0.21)	1.50 (0.20)	2.12 (0.01)	1.01 (0.03)	0.91 (0.02)**

Table 3: Precipitation and Temperature: Average and standard deviation RMSE for all scenarios. HMTL presented the best results. Two-sampled t-test was performed and statistically significant differences between HMTL and the second best method at a level of $p \leq 0.05$ (*) and $p \leq 0.01$ (**) are highlighted.

able prediction capacity in situations with limited amount of measurements (samples), as the results shown in Table 2.

Concluding Remarks

A hierarchical multitask learning (HMTL) framework to deal with multiple MTL problems is proposed. It was motivated by the problem of constructing Earth System Models (ESMs) ensemble for the simultaneous prediction of multiple climate variables. The formulation allows two levels of information sharing: (1) model parameters (coefficients of linear regression); and (2) precision matrices, which encodes the relationship of linear regressors. A group lasso regularization is responsible for capturing similar sparsity patterns across multiple precision matrices.

Experiments on joint projection of temperature and precipitation in South America showed that the HMTL produced more accurate predictions in many situations, when compared to the independent execution of existing MTL methods for each climate variable. Simulations on synthetic datasets also showed that the proposed HMTL achieved higher performance as the number of internal MTL problems increase. Here, only temperature and precipitation were used, as they are two of the most studied variables in the climate literature. Future works include a wider analysis with other climate variables, such as geopotential heights and wind directions.

Acknowledgments

We thank the anonymous reviewers for their valuable comments. AB was supported by NSF grants IIS-1563950, IIS-1447566, IIS-1447574, IIS-1422557, CCF-1451986, CNS-1314560, IIS-0953274, IIS-1029711, NASA grant NNX12AQ39A, and gifts from Adobe, IBM, and Yahoo.

FJVZ thanks CNPq for the financial support (Process no. 309115/2014-0).

References

- Danaher, P.; Wang, P.; and Witten, D. M. 2014. The joint graphical lasso for inverse covariance estimation across multiple classes. *J. R. Stat. Soc. Series B Stat. Methodol.* 76(2):373–397.
- Friedman, J.; Hastie, T.; and Tibshirani, R. 2008. Sparse inverse covariance estimation with the graphical lasso. *Bio-statistics* 9 3:432–441.
- Gonçalves, A. R.; Das, P.; Chatterjee, S.; Sivakumar, V.; Von Zuben, F. J.; and Banerjee, A. 2014. Multi-task Sparse Structure Learning. In *ACM CIKM*, 451–460.
- Gonçalves, A. R.; Von Zuben, F. J.; and Banerjee, A. 2016. Multi-task sparse structure learning with gaussian copula models. *JMLR* 17(33):1–30.
- Honorio, J., and Samaras, D. 2010. Multi-task learning of Gaussian graphical models. In *ICML*, 447–454.
- IPCC. 2013. Intergovernmental Panel on Climate Change Fifth Assessment Report.
- Kang, Z.; Grauman, K.; and Sha, F. 2011. Learning with whom to share in multi-task feature learning. In *ICML*, 521–528.
- Kumar, A., and Daume III, H. 2012. Learning task grouping and overlap in multi-task learning. In *ICML*, 1383–1390.
- Liu, D. C., and Nocedal, J. 1989. On the limited memory BFGS method for large scale optimization. *Mathematical Programming* 45(1-3):503–528.
- McQuade, S., and Monteleoni, C. 2013. MRF-Based Spatial Expert Tracking of the Multi-Model Ensemble. In *International Workshop on Climate Informatics*.

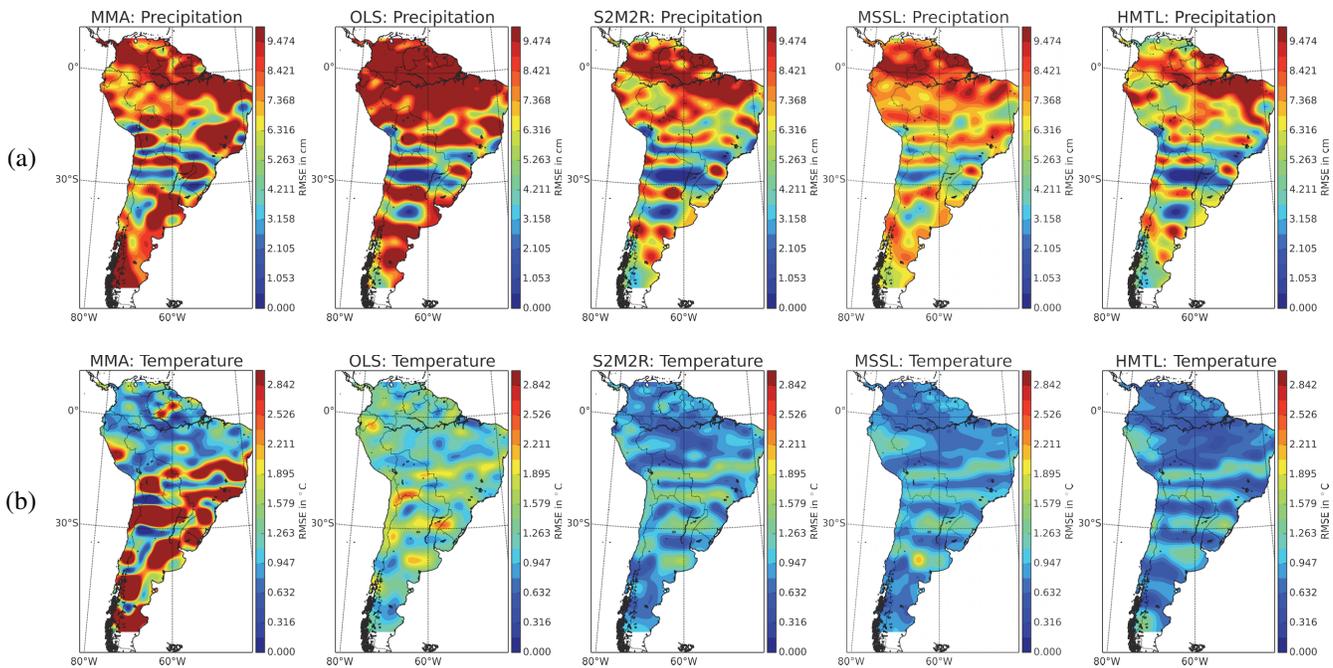


Figure 2: Precipitation - row (a), and Temperature - row (b), in summer (20 years): RMSE per geographical location for HMTL and four baselines. HMTL produced more accurate projections (lower RMSE) than the contenders.

Meinshausen, N., and Buhlmann, P. 2006. High-dimensional graphs and variable selection with the lasso. *Annals of Statistics* 34:1436–1462.

Mohan, K.; London, P.; Fazel, M.; Witten, D.; and Lee, S.-I. 2014. Node-based learning of multiple Gaussian graphical models. *JMLR* 15(1):445–488.

Sanderson, B. M.; Knutti, R.; and Caldwell, P. 2015. Addressing interdependency in a multimodel ensemble by interpolation of model properties. *J. of Climate* 28(13):5150–5170.

Subbian, K., and Banerjee, A. 2013. Climate Multi-model Regression using Spatial Smoothing. In *SIAM SDM*, 324–332.

Taylor, K.; Stouffer, R.; and Meehl, G. 2012. An overview of CMIP5 and the experiment design. *Bull. of the Am. Met. Soc.* 93(4):485.

Tebaldi, C., and Knutti, R. 2007. The use of the multi-model ensemble in probabilistic climate projections. *Phil. Trans. R. Soc. A* 365(1857):2053–2075.

Willmott, C. J., and Matsuura, K. 2001. Terrestrial Air Temperature and Precipitation: Monthly and Annual Time Series (1950 - 1999). http://climate.geog.udel.edu/~climate/html_pages/README_ghcn_ts2.html.

Yang, M.; Li, Y.; and Zhang, Z. 2013. Multi-task learning with Gaussian matrix generalized inverse Gaussian model. In *ICML*, 423–431.

Yuan, M., and Lin, Y. 2006. Model selection and estimation in regression with grouped variables. *J. R. Stat. Soc. Series B Stat. Methodol.* 68(1):49–67.

Zhang, Y., and Schneider, J. 2010. Learning multiple tasks with sparse matrix-normal penalty. In *NIPS*, 2550–2558.

Zhang, Y., and Yeung, D.-Y. 2010. A convex formulation for learning task relationships in multi-task learning. In *UAI*, 733–742.