Estimating Reference Evapotranspiration for Irrigation Management in the Texas High Plains

Daniel Holman
Texas Tech University
Lubbock, Texas, USA
deholman@ag.tamu.edu

Mohan Sridharan
Texas Tech University
Lubbock, Texas, USA
mohan.sridharan@ttu.edu

Prasanna Gowda
USDA-ARS
Bushland, Texas, USA
prasanna.gowda@ars.usda.gov

Dana Porter
Texas A&M AgriLife Extension Service
Lubbock, Texas, USA
d-porter@tamu.edu

Thomas Marek
Texas A&M AgriLife Research
Amarillo, Texas, USA
t-marek@tamu.edu

Terry Howell
USDA-ARS
Bushland, Texas, USA
terry.howell@ars.usda.gov

Jerry Moorhead
USDA-ARS
Bushland, Texas, USA
jed.moorhead@ars.usda.gov

Abstract

Accurate estimates of daily crop evapotranspiration (ET) are needed for efficient irrigation management in regions where crop water demand exceeds rainfall. Daily grass or alfalfa reference ET values and crop coefficients are widely used to estimate crop water demand. Inaccurate reference ET estimates can hence have a tremendous impact on irrigation costs and the demands on freshwater resources. ET networks calculate reference ET using precise measurements of meteorological data. These networks are typically characterized by gaps in spatial coverage and lack of sufficient funding, creating an immediate need for alternative sources that can fill data gaps without high costs. Although non-agricultural weather stations provide publicly accessible meteorological data, there are concerns that the data may be unsuitable for estimating reference ET due to factors such as weather station siting, data formats and quality control issues. The objective of our research is to enable the use of alternative data sources, adapting sophisticated machine learning algorithms such as Gaussian process models and neural networks to discover and model the nonlinear relationships between non-ET weather station data and the reference ET computed by ET networks. Using data from the Texas High Plains region in the U.S., we demonstrate significant improvement in estimation accuracy in comparison with baseline regression models typically used for irrigation management applications.

1 Introduction

Accurate estimates of daily crop evapotranspiration (ET) are essential for irrigation and water management within arid, semi-arid and semi-humid regions where crop water demand exceeds rainfall. A common method for estimating crop water demand uses daily grass or alfalfa reference ET (ET_o) values and crop coefficients (K_c). Accurate ET estimates are essential given the increased demands on U.S. freshwater resources, especially within the central Great Plains underlain by the vast but declining Ogallala aquifer. ET networks, such as the Texas High Plains ET Network (TXHPET), are collections of strategically located weather stations that gather meteorological data on well-watered reference crop (grass or alfalfa) to calculate reference ET. TXHPET consists of 15 weather stations covering a 39-county area of the Texas High Plains. Most of the region is semi-arid with temporal and spatial variability in precipitation. The region has high evaporative demand ($\approx$ 2500 mm/year Class A pan evaporation) due to high solar radiation, high vapor pressure deficit (VPD) and strong regional advection. Due to limitations in spatial coverage of the 15 weather stations, it is a challenge to accurately estimate ET in the associated counties. Although a sufficiently dense network can capture the spatial variability of parameters used to compute reference ET, funding and staffing issues restrict ET networks from expanding coverage through additional weather stations. There is thus an immediate need for exploring alternative data sources capable of augmenting data gaps without high maintenance and field-based support costs. While data from non-agricultural, non-ET weather stations are publicly accessible, there are concerns that the data may not be appropriate for estimating reference ET due to factors such as weather station siting, data formats, fetch requirements and data quality issues.

The research reported in this paper seeks to identify, evaluate and use alternative meteorological data sources for augmenting ET networks. As a representative example, we consider the TXHPET network and a publicly available non-ET network maintained by the National Weather Service (NWS), henceforth referred to as "*ET stations*" and "*NWS stations*" respectively. Sophisticated machine learning algorithms such as Gaussian process models and neural networks are used to formulate the challenge of discovering and modeling the nonlinear relationships between meteorological data from NWS stations and reference ET computed by ET stations; the learned models are then used to estimate reference ET based on data from NWS stations. We compare the estimation accuracy of these algorithms with that of linear regression models typically used for irrigation management applications.

The research described in this paper has broader implications for water sustainability. The experimental methodology can be used for reference ET computation from non-agricultural and non-ET weather stations in other regions of the world. Accurate reference ET estimates will help producers accurately estimate crop water use and manage their soil-water profiles, preventing over or under-watering of crops to minimize crop damage and loss of profits. Furthermore, the use of sophisticated machine learning algorithms, which are yet to be fully exploited for water management, can impact conservation of groundwater resources.

2 Related Work

Reference ET can be calculated using either the FAO-56 [Allen *et al.*, 1998] equation or the ASCE Standardized Reference ET Equation [Allen *et al.*, 2005]. Areal coverage is not universal and there are significant gaps in the spatial coverage. It is further complicated by high spatial variations in air temperature, wind speed, wind direction and other weather parameters due to regional effects such as atmospheric circulation patterns and local effects such as topography, land use, elevation and soil properties. It is hence difficult to determine daily reference ET for irrigation management of large regions using one (or a small set of) predetermined station(s). Although remote sensing-based ET estimates are showing promise of expanding areal coverage and integration capabilities, accuracy depends on accuracy of input weather data.

Most statistical models reported in irrigation management literature are based on ordinary least square regression. Popular models used in regression include: (1) linear; $Y = a + bX$; (2) exponential; $Y = ae^{bX}$; (3) power or logarithmic; $Y = aX^b$; and (4) a quadratic polynomial; $Y = aX^2 + bX + c$. Variable Y represents the desired output (e.g., reference ET values from ET stations) and X represents input values such as rainfall, irrigation amount, weather parameters or reference ET estimated at NWS stations. Values of coefficients a, b, and c are tuned on training data such that computed values of output are as close as possible to the given (i.e., ground truth) output values. These regression formulations tend to fix the basis functions before observing training data, and the number of basis functions grows exponentially with the dimensions of input space. Furthermore, the basis functions are not adaptable to data and the curse of dimensionality makes a strong case for more sophisticated models.

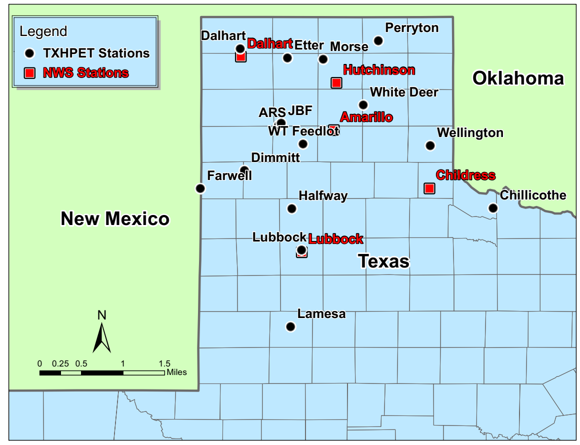

Figure 1: Locations of 15 ET network stations and five NWS weather stations in the Texas High Plains.

The machine learning (ML) research community has developed many sophisticated algorithms for learning, inference, and estimation. In recent years, there has been a push towards using ML algorithms to address concrete real-world problems [Wagstaff, 2012]. Popular ML algorithms such artificial neural networks (ANNs), support vector machines (SVMs), and Gaussian Process models (GPs) provide substantial benefits over linear regression models [Bishop, 2008]. For instance, ANNs can be used to adaptively model complex functions between input and output parameters, while SVMs project input features to high dimensions, resulting in sparse representations and robust decision boundaries. Similarly, GPs use a non-parametric kernel-based algorithm to capture the evolution of normally-distributed random variables representing the patterns being tracked [Higdon *et al.*, 2003].

Some ML algorithms have been used for estimation tasks in the realm of water sustainability. For instance, ANNs have been used to design simulation and forecasting models for rainfall-runoff [Campolo *et al.*, 1999; Hsu *et al.*, 1995; Tokar and Johnson, 1999; Zealand *et al.*, 1999], remote sensing applications [Sudheer *et al.*, 2010], and for downscaling in remote sensing-based irrigation management [Ha *et al.*, 2011]. Similarly, GPs have been used to model wave characteristics in oceanographic data [Rychlik *et al.*, 1997]. However, complex irrigation management challenges have not yet been formulated using sophisticated machine learning algorithms. This paper investigates the use of GPs and ANNs, which are well-suited to discover and model the nonlinear relationships involved in reference ET estimation.

3 Materials and Methods

In an elaborate initial data preparation phase, publicly accessible weather networks in the Texas High plains were assessed based on real-time availability of data and continuity of historic records of necessary parameters. Five suitable NWS (non-ET) stations were identified and their locations were used to create a Thiessen polygon map, which (in turn) was used to pair these stations with appropriate ET stations of the TXHPET network for evaluation—Figure 1 shows the ET stations and NWS stations. The relevant weather parameters include: minimum and maximum daily air temperature, dew point temperature, relative humidity, solar radiation, wind speed, and barometric pressure; missing data were estimated using the standardized reference ET methodology [Allen *et al.*, 2005]. Models were then learned, e.g., using GP and regression algorithms, to understand the relationships between each weather parameter measured at NWS stations and the corresponding measurements at ET stations. Certain parameters are expected to have lower correlation, due to station siting and spatial variability (e.g. wind speed) or due to manual computation (e.g. solar radiation and barometric pressure). Models were also learned using different subsets of parameters measured at NWS stations as input and the reference ET computed at ET stations as output. All models were evaluated using coefficient of determination ($R^2 \in [0, 1]$) and root mean square error (RMSE). R^2 describes the proportion of variability in observed data that is explained by the model; a higher value indicates a better fit—see Moriasi et al. (2007) for more information on performance statistics. These models helped confirm that reference ET computation requires models of nonlinear relationships and needs to consider all NWS weather parameters—more details in Section 4.

Reference ET estimation from weather parameters measured by NWS stations was then posed as a supervised learning problem. The input is a set of vectors of seven weather parameters collected by NWS weather stations and used to compute reference ET. The output (i.e., target) values are reference ET values from the TXHPET network. Learned models capture the relationship between inputs and output, and estimate reference ET given new data from the NWS stations. Specifically, training data consists of N vectors of inputs: $\mathbf{X} = \{\mathbf{x_1}, ..., \mathbf{x_N}\}$ and corresponding target outputs: $\mathbf{T} = \{t_1, ..., t_N\}$, while testing data consists of previously unseen inputs and target outputs: $\{(\hat{\mathbf{x}}_\mathbf{i}, \hat{t}_i), i = 1...N\}$. Learned models process input vectors of the testing dataset to estimate outputs: $\hat{y}(\mathbf{x})$, which are compared with ground truth target outputs. Models were also learned using subsets of input vectors to determine if certain parameters were not relevant to reference ET computation. This paper investigates the use of Gaussian Process models (GPs) and Artificial Neural Networks (ANNs), and compares their estimation accuracy with (baseline) linear regression (LR) models using performance statistics (R^2 and $RMSE$).

The experimental evaluation used the daily reference ET database over a period of ten years (2001-2010), which was divided into two equal parts. Data corresponding to odd-numbered days of the year were used for model development and data from even-numbered days of the year were used for validation. Such a data division is required because the wet and dry years are typically inconsistent; estimation models trained with data corresponding to specific years (or months) may result in high estimation errors on data corresponding to other years (or months). Our data division ensures that test data are drawn from the same population as the training data. To eliminate any bias, experiments were repeated after swapping the training and test datasets.

3.1 Linear Regression

Linear regression (LR) models constitute one of the simplest forms of regression algorithms. LR models are widely used for estimation problems in the irrigation management and water sustainability community; we used LR models as the baseline for comparison. Ordinary least squares regression is the most common form of LR. Estimating output values for input vectors is posed as the task of evaluating parameters of a linear equation of the form:

$$y(\mathbf{x}, \mathbf{w}) = w_0 + w_1 x_1 + ... + w_D x_D \tag{1}$$

where y is the dependent (output or target) variable that is a linear function of the input vector $\mathbf{x}$ and weights $\mathbf{w}$; and w_0 is the *bias* parameter. For computing reference ET (described in Section 3),, $\mathbf{x}$ is the vector of (NWS) weather parameter values and y is the target reference ET value. Each entry x_i of the input vector $\mathbf{x}$ has a corresponding w_i that needs to be determined to best fit the training data. Although it is possible to include nonlinear *basis functions* (see below) of input variables in LR, the simple LR models widely used in irrigation management applications form the baseline. LR models (in general) have useful computational properties, but they use fixed basis functions and are limited in their applicability to large scale problems by the curse of dimensionality.

3.2 Artificial Neural Network

Artificial neural networks (ANNs) were motivated by models created to represent information processing in biological systems [Bishop, 2008]; the human brain is known to contain many billion neurons that are interconnected to form a network. Within the ML community, ANNs are used as models for statistical pattern recognition; nodes of the network represent inputs (and hidden variables) and the edges represent the influence nodes have on each other. ANNs build on the more advanced linear models of regression that use linear combinations of *fixed* nonlinear basis functions $\phi_j(\mathbf{x})$:

$$y(\mathbf{x}, \mathbf{w}) = f\Big(\sum_{j=1}^{M} w_j \phi_j(\mathbf{x})\Big) \tag{2}$$

ANNs make this model adaptive by parameterizing the basis functions. Typically each basis function is a nonlinear function of a linear combination of inputs; the coefficients of this linear combination and the w_j in Equation 2 are parameters that are tuned adaptively to best fit training data.

A popular type of ANN is the multilayer feed forward network in which data flow from input to output. A feed forward ANN can have multiple layers; the first (input) layer connects to input variables, and the final (output) layer connects to target variables. One or more *hidden* layers can lie between input and output layers. The processing elements (i.e., nodes) of each layer connect with neighboring layers; typically, nodes within the same layer are not connected and nodes in non-adjacent layers are not connected. Each node in a layer (other than input layer) typically receives a signal from all nodes of the previous layer. The effective input signal at each node is the weighted sum of previous layer nodes' outputs. This input signal passes through a nonlinear function to produce the output signal of the node. The ANN thus represents a set of functional transformations [Bishop, 2008]. For an ANN with one hidden layer, *activations* are first constructed as linear combinations of input variables:

$$a_j = \sum_{i=1}^{D} w_{ji}^{(1)} x_i + w_{j0}^{(1)} \tag{3}$$

where $j = 1, ..., M$, and superscript (1) indicates that the parameters correspond to first layer of the network. These activations are the inputs of hidden layer nodes, and are transformed using nonlinear *activation functions*: $z_j = h(a_j)$, providing outputs of the hidden units, i.e., outputs of basis functions in Equation 2. Typically, the logistic sigmoid function is used for this nonlinear transformation: $h(a) = 1/(1 + exp(-a))$. Outputs of these transformations are combined linearly to provide *output unit activations*:

$$a_k = \sum_{j=1}^{M} w_{kj}^{(2)} z_j + w_{k0}^{(2)} \tag{4}$$

where $k = 1, ..., K$, with K being the total number of output nodes, and superscript (2) indicates that the parameters correspond to second layer of the network. The output activations are transformed using the appropriate activation function (identity function for standard regression) to obtain network outputs y_k. The overall network function for sigmoidal output unit activation functions combines Equations 2–4:

$$y_k(\mathbf{x}, \mathbf{w}) = \sigma \left(\sum_{j=1}^{M} w_{kj}^{(2)} h \left(\sum_{i=1}^{D} w_{ji}^{(1)} x_i + w_{j0}^{(1)} \right) + w_{k0}^{(2)} \right) \tag{5}$$

where vector $\mathbf{w}$ includes weights and bias parameters, controlling the ANN that is a nonlinear mapping from inputs to outputs. For reference ET computation, the network has seven input nodes and a single output node ($K = 1$). During the training phase, the *backpropagation* algorithm repeatedly performs a forward pass of data (through the network) to compute the error between actual output and network output, and performs a backward pass of the gradient of error to change $\mathbf{w}$ to best fit the data.

3.3 Gaussian Process

Gaussian processes are sophisticated learning models used for regression [Williams and Rasmussen, 1996] and classification [Williams and Barber, 1998]. Parametric models such as ANNs and linear regression require $y(\mathbf{x})$ to have an explicitly defined functional form whose parameters are defined in advance. The Bayesian treatment of regression starts with a prior distribution over weights $\mathbf{w}$ and $y(\mathbf{x}, \mathbf{w})$, and uses training data to obtain a posterior distribution over regression functions that is used for estimation. In GP models, the parameters are eliminated and the unknown function $y(\mathbf{x})$ exists in the infinite-dimensional space of possible functions, and a prior is defined over this space. GP models are thus considered to be non-parametric. The difficulty of revising a distribution in the infinite space of functions is overcome by considering values of the function over the discrete set of training samples [Bishop, 2008]. Stochastic random variables define priors for each input vector, and random functions defined over the space of inputs constitute the GP prior. During the training phase, the discrete set of inputs are used to modify these functions to pass as close as possible to the target outputs, thus approximating the (unknown) underlying function. Gaussian processes can be viewed as a natural generalization of a Gaussian distribution over a finite vector space to an infinite space of functions. A GP is given by:

$$f \sim GP(\mu(x), C(x, x')) \tag{6}$$

where the function f is distributed as a Gaussian process with mean function $\mu(x)$ and covariance function $C(x, x')$. In this paper,, the mean function is defined as the zero function. The covariance function expresses the expected covariance of values at each pair of points $\mathbf{x}$ and $\mathbf{x}'$. Given N input vectors in the training data, the covariance function is a $N \times N$ matrix $K : K_{ij} = C(\mathbf{x_i}, \mathbf{x_j})$. This matrix can be used to estimate (output) values for new inputs. In general, the estimated distribution is Gaussian with mean and covariance:

$$\hat{y} = \mathbf{k}^T(\mathbf{x}) K^{-1} \mathbf{t} \tag{7}$$
$$\sigma_{\hat{y}}^2(\mathbf{x}) = C(\mathbf{x}, \mathbf{x}) - \mathbf{k}^T(\mathbf{x}) K^{-1} \mathbf{k}(\mathbf{x})$$

where $\mathbf{k}(\mathbf{x}) = (C(\mathbf{x}, \mathbf{x}_1), ..., C(\mathbf{x}, \mathbf{x}_N))^T$, K is the covariance matrix for training data, and $\mathbf{t} = (t_1, ..., t_N)^T$. The default algorithm for training GP models has $O(N^3)$ time complexity due to matrix inversion in Equation 7. GP formulations can therefore become infeasible for data with a very large number of samples. Although algorithms are being developed to enable GP formulations of datasets that have a large number of samples, time complexity is not an issue of concern for the experiments reported in this paper.

Many different options exist for selecting covariance functions for a GP. The main requirement is that the function should generate a non-negative definite covariance matrix for any set of inputs $(\mathbf{x}_1, ..., \mathbf{x}_N)$. Graphically, the goal is to define covariances such that points that are nearby in the input space produce similar estimations. The research reported in this paper uses the popular radial basis function (RBF) kernels [Musavi *et al.*, 1992]:

$$C(x, x') = e^{-\gamma * (x - x')^2} \tag{8}$$

GP models do not require the tuning of specific parameters or weights. Instead, the covariance function contains hyper-parameters that can be tuned automatically to maximize the likelihood of training data. Assigning different values to hyper-parameters results in different GP models. We randomly initialized a finite set of hyper-parameter values over the space of possible values. The *training error* is calculated using each such GP model learned from training data inputs, comparing estimated outputs with actual outputs in the training data. Hyper-parameter values that provide GP models with the lowest training error are selected.

4 Experimental Results

The GP, ANN, and LR models were evaluated on the ten-year historical data for the Texas High Plains region; five NWS stations were paired with an appropriate subset of 15 TXHPET stations. Models were learned using 50% of the available data

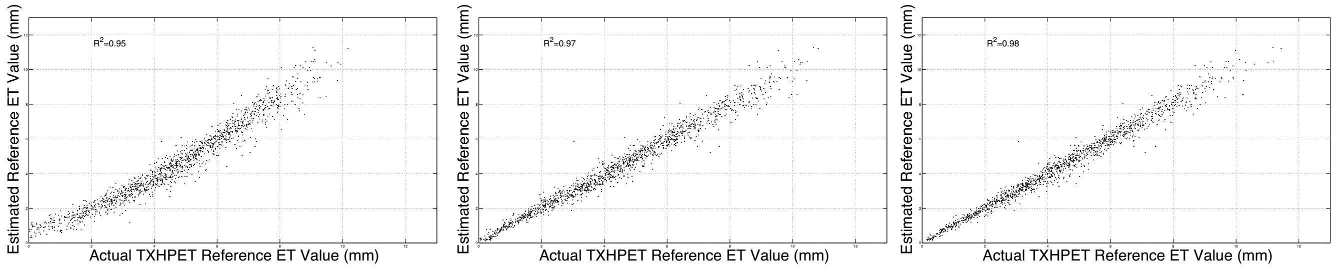

(a) LR Estimations (b) ANN Estimations (c) GP Estimations

Figure 2: Illustrative example of the estimation capabilities of LR, ANN, and GP models on the data obtained from one NWS station-ET station matched pair.

Table 1: Performance measures for estimations obtained from the LR, ANN, and GP models at each NWS station-ET station pair considered in this study. GP models provide highest accuracy while ANN models are more accurate than LR models.

	Linear Regression		Artificial Neural Network		Gaussian Process	
NWS - TXHPET Station	R^2	$RMSE(mm)$	R^2	$RMSE(mm)$	R^2	$RMSE(mm)$
Amarillo - Bushland-ARS	0.90	0.80	0.94	0.63	0.95	0.60
Amarillo - Dimmit	0.90	0.80	0.92	0.69	0.92	0.68
Amarillo - Bushland-JBF	0.90	0.85	0.94	0.65	0.95	0.62
Amarillo - West Texas A&M Feedlot	0.90	0.80	0.94	0.62	0.95	0.58
Childress - Chillicothe	0.87	0.93	0.90	0.80	0.91	0.76
Childress - Wellington	0.88	0.85	0.92	0.74	0.92	0.70
Dalhart - Dalhart	0.95	0.54	0.98	0.36	0.98	0.33
Dalhart - Etter	0.92	0.74	0.94	0.62	0.95	0.59
Hutchinson - Morse	0.90	0.84	0.94	0.65	0.95	0.61
Hutchinson - Perryton	0.88	0.94	0.93	0.72	0.94	0.69
Hutchinson - White Deer	0.90	0.83	0.94	0.64	0.94	0.62
Lubbock - Farwell	0.87	0.84	0.88	0.79	0.89	0.76
Lubbock - Halfway	0.91	0.71	0.94	0.59	0.94	0.57
Lubbock - Lamesa	0.91	0.73	0.93	0.62	0.94	0.58
Lubbock - Lubbock	0.94	0.58	0.97	0.4	0.98	0.36

and evaluated on the remaining 50% (previously unseen) of the data—see Section 3. The algorithms were implemented in Java by building on the WEKA open-source machine learning library [Hall *et al.*, 2009].

Models learned to discover the relationships between NWS station parameters and the corresponding ET station parameters show high correlation in certain parameters (e.g., $R^2 = 0.99$ for air temperature), while parameters with high spatial variability, e.g. wind speed and solar radiation, have a much lower correlation. These observations may partly be due to siting of NWS stations, e.g., micro-climate and proximity to fields or roads; much lower correlations were found when spatial distance between NWS and ET stations were larger. Accuracy of the models also decreased with increase in the spatial distance, which may be because the stations fall under different atmospheric circulation patterns [Buishand and Brandsma, 1997]. Models learned using subsets of weather parameters measured at NWS and the reference ET computed at ET stations identified the presence of nonlinear relationships and confirmed that all NWS weather parameters need to be considered for reference ET computation. All subsequent experiments (below) therefore used all NWS weather parameters as inputs and the reference ET computed at ET stations as output to learn and evaluate models.

Figure 2 illustrates the estimation capabilities of LR, ANN, and GP model for a specific NWS station-ET station matched pair. The estimation capability is evaluated using the R^2 measure; models that provide highly accurate estimations will result in points that lie on (or very close to) the $Y = X$ line. These plots show that the estimations provided by ANN and GP models are significantly more accurate than the LR models predominantly used for irrigation management applications. We also observe that GP models perform slightly better than ANN models—most points are along the diagonal in Figure 2(c) and the R^2 value is larger.

Table 1 summarizes the performance of each of the three models at each NWS station-ET station pair considered in our study. The estimation capability was evaluated using the R^2 and $RMSE$ performance measures. Similar to the results observed in Figure 2, the ANN and GP models provide more accurate estimations as they better capture the relationships in the data. We also observed that the GP models performed at least as well as the the ANN models in all NWS-ET station pairs considered in the experimental trials; the difference is more pronounced in certain instances and results are statistically significant. With the GP models, the R^2 values are closer to 1 with a lower $RMSE$.

Figure 3 shows the RMSE measures obtained for the 15

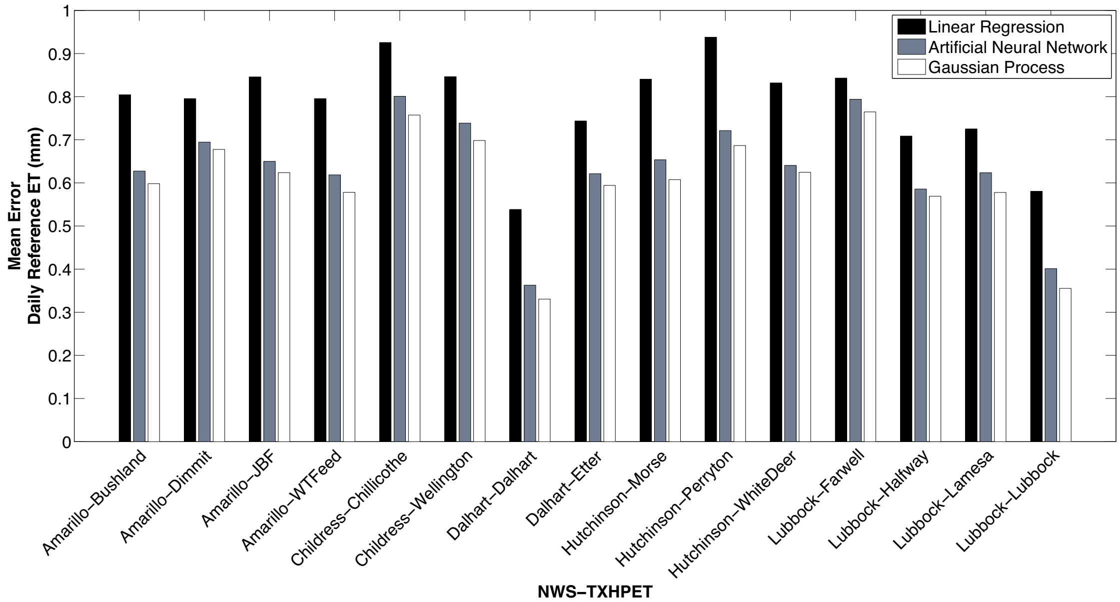

Figure 3: Comparison of RMSE obtained using the GP, ANN, and LR models for each NWS station-ET station pair considered in this study. GP and ANN models provide much lower RMSE in comparison with LR models, with the GP models performing slightly better than the ANN models.

NWS station-ET station pairs. GP and ANN models provide much lower RMSE in comparison with LR models, i.e., reference ET values estimated by GP and ANN models are much closer to reference ET values obtained from the corresponding TXHPET stations. Although the improvement in estimation accuracy of the GP and ANN models (compared with LR models) is different at different stations, the improvement is statistically significant for all stations, including the improvements provided by GP in comparison with ANN. GP models thus show significant promise in enabling the use of alternative data sources for computing reference ET values.

Pairing the NWS stations with the corresponding ET stations at Lubbock and Dalhart produced highly accurate estimations: $R^2 = 0.98$ and 0.98 respectively. However, matching the NWS station at Lubbock with the Farwell ET station obtained $R^2 = 0.89$ and RMSE of 0.76 mm for daily reference ET.; although this represents $\approx 29\%$ error, it is still significantly better than LR models that provide an RMSE of 0.84 mm with a relative error of 33%. Future research will consider additional features for input (e.g., elevation) and other GIS selection methods for pairing NWS stations with TXHPET stations. For instance, although our analysis identifies a good correlation between the Farwell TXHPET station and the Lubbock NWS station, including additional features may help identify stations that are strongly correlated.

The improvement in estimation accuracy provided by ANN and GP models has significant practical value. For instance, the average improvement in RMSE for daily reference ET estimations provided by GP and ANN models compared with the LR models is ≥ 0.2 mm. For a typical crop season of about 200 days, this amounts to approximately 40 mm over the season. Although this difference may seem rather small, one acre-inch of water for all fields of the Texas High Plains translates to ≈ 93.9 billion liters of water [Marek *et al.*, 2010], which is comparable to the amount of water used by the entire city of Houston, Texas (population ≈ 2 million in 2010 US census) in about two and a half months.

5 Conclusion

Efficient water resource management is a pressing concern in irrigated agriculture throughout arid and semi-arid regions in which crop water demands exceed rainfall. Accurate estimates of crop evapotranspiration (ET) are essential for irrigation management in these regions. Due to spatial variability and funding constraints for dense networks, ET stations struggle to accurately estimate reference ET in their entire areal coverage. This paper describes the results of a research study towards the long-term objective of using publicly available non-agricultural, non-ET stations for filling the data gaps in ET networks.

We investigated the use of sophisticated machine learning algorithms such as artificial neural networks (ANNs) and Gaussian process models (GPs) to discover and model the nonlinear relationships between meteorological data collected by National Weather Service (NWS) stations and the reference ET computed by ET stations in the Texas High Plains. Experimental results show that ANNs and GPs provide significantly more accurate estimations of daily reference ET than the linear regression models widely used for irrigation management applications; the (non-parametric) GP models result in more accurate estimates of reference ET in comparison with ANNs. This improvement in estimating the reference ET values translates to huge reductions in costs associated with wasteful use of water resources (due to overwatering), in addition to minimizing crop stress and crop loss (due to under-watering).

Although this study focused on the Texas High Plains, the experimental methodology can be adapted to other regions of the world. Furthermore, the machine learning algorithms described in this paper in the context of a key irrigation management challenge possess significant potential for addressing open challenges in water resources management and other subfields of agriculture.

References

[Allen *et al.*, 1998] Richard G. Allen, Luis S. Pereira, Dirk Raes, and Martin Smith. Crop Evapotranspiration - Guidelines for Computing Crop Water Requirements. Technical Report 56, Food and Agriculture Organization of the United Nations, ROME, 1998.

[Allen *et al.*, 2005] Richard G Allen, Ivan A. Walter, Ronald Elliot, and Terry A. Howell. *ASCE-EWRI Standardization of Reference Evapotranspiration*. American Society of Civil Engineers - Environmental Water Institute, 2005.

[Bishop, 2008] Christopher M. Bishop. *Pattern Recognition and Machine Learning*. Springer-Verlag, New York, 2008.

[Buishand and Brandsma, 1997] T. A. Buishand and T. Brandsma. Comparison of Circulation Classification Schemes for Predicting Temperature and Precipitation in the Netherlands. *International Journal of Climatology*, 17:875–889, 1997.

[Campolo *et al.*, 1999] Marina Campolo, Paolo Andreussi, and Alfredo Soldati. River flood forecasting with a neural network model. *Water resources research*, 35(4):1191–1197, 1999.

[Ha *et al.*, 2011] Wonsook Ha, Prasanna H. Gowda, and Terry A. Howell. A review of downscaling methods for remote sensing-based irrigation management: part i. *Irrigation Science*, pages 1–20, 2011.

[Hall *et al.*, 2009] M. Hall, E. Frank, G. Holmes, B. Pfahringer, P. Reutemann, and I.H. Witten. The weka data mining software: an update. *ACM SIGKDD Explorations Newsletter*, 11(1):10–18, 2009.

[Higdon *et al.*, 2003] David M. Higdon, Herbert KH Lee, and Christopher H. Holloman. Markov Chain Monte Carlo-based Approaches for Inference in Computationally Intensive Inverse Problems. In J. M. Bernardo, M. J. Bayarri, J. O. Berger, and A. P. Dawid, editors, *Bayesian Statistics*, volume 7, pages 181–197. Oxford University Press, 2003.

[Hsu *et al.*, 1995] Kuo-Lin Hsu, Hoshin Vijai Gupta, and Soroosh Sorooshian. Artificial neural network modeling of the rainfall-runoff process. *Water resources research*, 31(10):2517–2530, 1995.

[Marek *et al.*, 2010] Thomas H. Marek, Dana O. Porter, Prasanna H. Gowda, Terry A. Howell, and Jed E. Moorhead. Assessment of Texas Evapotranspiration (ET) Networks, 2010.

[Moriasi *et al.*, 2007] Daniel N. Moriasi, Jeffrey G. Arnold, Michael W. Van Liew, Ronald L. Bingner, R. Daren Harmel, and Tamie L. Veith. Model Evaluation Guidelines for Systematic Quantification of Accuracy in Watershed Simulations, 2007.

[Musavi *et al.*, 1992] Mohamad T. Musavi, Wahid Ahmed, Khue Hiang Chan, Kathleen B. Faris, and Donald M. Hummels. On the training of radial basis function classifiers. *Neural networks*, 5(4):595–603, 1992.

[Rychlik *et al.*, 1997] Igor. Rychlik, Pär Johannesson, and M. R. Leadbetter. Modeling and statistical analysis of ocean-wave data using transformed gaussian processes. *Marine Structures*, 10(1):13–47, 1997.

[Sudheer *et al.*, 2010] K. P. Sudheer, Prasanna H. Gowda, Indrajeet Chaubey, and Terry A. Howell. Artificial neural network approach for mapping contrasting tillage practices. *Remote Sensing*, 2(2):579–590, 2010.

[Tokar and Johnson, 1999] A. Sezin Tokar and Peggy A. Johnson. Rainfall-runoff modeling using artificial neural networks. *Journal of Hydrologic Engineering*, 4(3):232–239, 1999.

[Wagstaff, 2012] Kiri Wagstaff. Machine Learning that Matters. In *International Conference on Machine Learning*, Edinburgh, Scotland, June 25 - July 1 2012.

[Williams and Barber, 1998] Christopher KI Williams and David Barber. Bayesian classification with gaussian processes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 20(12):1342–1351, 1998.

[Williams and Rasmussen, 1996] Christopher KI Williams and Carl Edward Rasmussen. Gaussian processes for regression. *Advances in Neural Information Processing Systems*, 1996.

[Zealand *et al.*, 1999] Cameron M. Zealand, Donald H. Burn, and Slobodan P. Simonovic. Short term streamflow forecasting using artificial neural networks. *Journal of hydrology*, 214(1):32–48, 1999.