

Attention-Based Models for Snow-Water Equivalent Prediction

Krishu K Thapa¹, Bhupinderjeet Singh¹, Supriya Savalkar¹, Alan Fern², Kirti Rajagopalan¹, Ananth Kalyanaraman¹

¹Washington State University, Pullman, WA

²Oregon State University, Corvallis, OR

{krishu.thapa, bhupinderjeet.singh, supriya.savalkar}@wsu.edu, alan.fern@oregonstate.edu, {kirtir, ananth}@wsu.edu

Abstract

Snow Water-Equivalent (SWE)—the amount of water available if snowpack is melted—is a key decision variable used by water management agencies to make irrigation, flood control, power generation and drought management decisions. SWE values vary spatiotemporally—affected by weather, topography and other environmental factors. While daily SWE can be measured by Snow Telemetry (SNOTEL) stations with requisite instrumentation, such stations are spatially sparse requiring interpolation techniques to create spatiotemporally complete data. While recent efforts have explored machine learning (ML) for SWE prediction, a number of recent ML advances have yet to be considered. The main contribution of this paper is to explore one such ML advance, attention mechanisms, for SWE prediction. Our hypothesis is that attention has a unique ability to capture and exploit correlations that may exist across locations or the temporal spectrum (or both). We present a generic attention-based modeling framework for SWE prediction and adapt it to capture spatial attention and temporal attention. Our experimental results on 323 SNOTEL stations in the Western U.S. demonstrate that our attention-based models outperform other machine learning approaches. We also provide key results highlighting the differences between spatial and temporal attention in this context and a roadmap toward deployment for generating spatially-complete SWE maps.

Introduction

Much of the high-value, productive irrigated agriculture in the Western U.S. is facilitated by snowmelt-driven streamflow. Many of these watersheds are “fully appropriated” — that is, every drop of water is appropriated to a particular use under a legal water rights system. Irrigation needs to be balanced with other uses such as hydropower production, municipal needs, and environmental flows. In managing these diverse uses in water-scarce environments, key decision variables used by water management agencies include an expectation of the snowpack—measured as *Snow Water-Equivalent (SWE)*—and the timing of snowmelt. SWE is the amount of water that would be available if the snowpack is melted. For example, the U.S. Bureau of Reclamation makes decisions on storage and release of water from reservoirs to

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

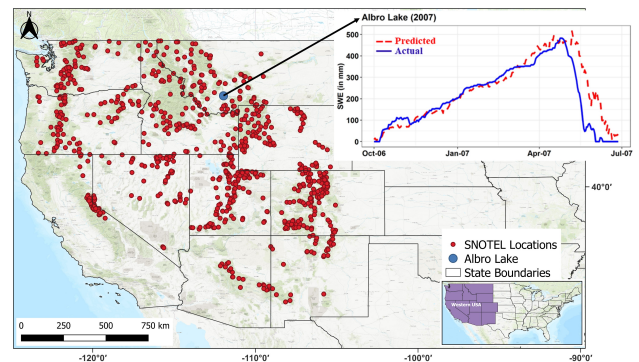


Figure 1: Spatial map showing the sparse spread of 822 SNOTEL locations across Western U.S., of which we used a subset of 323 stations based on data availability. The inset shows two SWE curves for a station (Albro Lake) for the 2007 SWE season. The blue curve is ground-truth (SNOTEL) and the red curve is the predicted curve obtained by our ensemble attention model.

meet diverse needs via forecasts of inflows into the reservoir, which uses SWE as a key input.

Given the importance of this information, over the last few decades, large investments were made in measurements via the Snow Telemetry (SNOTEL) station network across the Western U.S. These stations report the daily SWE value along with other related attributes including (but not limited to) snow depth, temperature, precipitation, and relative humidity. However, these stations are sparsely distributed in space (Figure 1). The 822 SNOTEL stations shown average to around one station for every 4,014 square kilometers.

Currently, long-term spatially-complete SWE products are developed via interpolation from SNOTEL and other observational networks (Broxton, Dawson, and Zeng 2016). (For a brief review of current tools, see Section Related Work). These methods do not account for large variations in SWE in short distances and take a simplistic approach to incorporating the effects of environmental variables.

While SNOTEL information may be spatially-incomplete, there are other types of data such as remotely-sensed reflectances (Brodzik et al. 2016) or topographical details, and weather networks (providing

temperature, precipitation etc.) that are denser than the SNOTEL network. Leveraging these data along with state-of-the-art AI models can significantly enhance our ability to obtain snowpack information at watershed scales and significantly improve decision-making. Such data can advance hydrological models by reducing the uncertainties involved in their parameterization of snow processes.

In this paper, we address the problem of predicting daily SWE for any location (with or without a SNOTEL station) on any given date, using observable data that are available at high spatial resolution (e.g., weather, satellite imagery). Our prediction problem is motivated by the above established need for generating spatially-complete snowpack products.

Once such a generic prediction model is designed, it can also be extended to forecasting scenarios.

Challenges: Predicting SWE for a particular location and time is a challenging task for multiple reasons. SWE is a complex trait affected by environmental variables and spatial attributes (e.g., elevation, latitude-longitude, topography, land cover) (Kim et al. 2021; Rutter et al. 2009; Sospedra-Alfonso, Melton, and Merryfield 2015). Time is another dimension affecting prediction. The window of time from when the snow starts to accumulate to when snow melts completely is referred to as the *SWE season*. While the SWE season varies geographically, an approximate 270-days window of (Oct. 1 to Jul. 1) covers the majority of the snow season for the Western U.S.

Figure 1 shows the general expected trend of the SWE curve over the SWE season for a SNOTEL station and year. However, there is significant heterogeneity in these curves across locations and seasons. For instance, the overall SWE values, maximum SWE and the date when the ground is snow-free, could all be different for a drought year versus a non-drought year.

Machine learning (ML) models have been recently explored for SWE prediction (Bair et al. 2018; Marofi, Tabari, and Abyaneh 2011; Meyal et al. 2020)—see Section Related Work. These approaches pose the problem as time series prediction and use deep learning models such as recurrent neural networks. However, they largely ignore spatial correlations.

Contributions: We explore *attention* as the main mechanism for SWE prediction. Self-attention was introduced as part of the transformer architecture (Vaswani et al. 2017). The main idea is to compute a set of attention weights for all token pairs in a sequence and use those weights to make predictions. Even though originally applied to machine translation, it has since been widely adopted in other applications that involve a sequence-to-sequence transformation.

We hypothesize that the attention mechanism can help the SWE prediction task by capturing and exploiting any correlations that may exist between locations (*spatial attention*) or between time intervals (*temporal attention*). For instance, in a study using all SNOTEL stations in the U.S., we observed strong correlations between locations that have a short spatial distance or similar elevation or both (Figure 2).

To test our working hypothesis, we first designed a generic attention-based model architecture and adapted it to




Figure 2: SWE correlation for all location pairs as a function of their spatial distance and elevation difference product.

perform spatial attention, temporal attention, and a simple ensemble version which benefits from both.

Our experiments on 323 SNOTEL locations in the Western U.S. demonstrate the general effectiveness of the attention based schemes, which outperform other prior approaches. Our study also highlights the differences between the spatial and temporal attention schemes. An extensive report on our experiments and findings are presented in the Section Experimental Results, followed by a roadmap (Section Roadmap to Deployment) toward deployment for generating a spatially-complete SWE map.

Related Work

Among the earlier work, Huang and Cressie (1996) developed a Kalman filter for predicting SWE. Their approach uses observed temperature and precipitation to simulate snow cover conditions. The Kalman filter assumes that the relationship between the independent and the dependent variables is linear. However, this assumption may or may not hold depending on the location and its various attributes including elevation and weather conditions. Other work used spatial interpolation techniques such as kriging (Carroll, Carroll, and Poston 1999) and statistical regression tree (Elder, Rosenthal, and Davis 1998). While these approaches use statistical models to interpolate by spatial proximity, they lack the sophistication to include additional features (e.g., weather, elevation, satellite reflectance observations) which could be key SWE indicators.

Examples of spatially-complete SWE estimates include Barrett (2003) and Broxton et al. (2016). The former is based on integration of numerical weather models and data assimilation but is only available starting 2003. The latter is based on simple interpolation of SNOTEL and other observational data and is available as a longer time series. However spatiotemporal correlations are not accounted for, and the effect of other environmental variables are considered in a simplistic manner. Neither explored ML methods. ML approaches have been explored more recently. Bair et al. (2018) used bagged trees and neural networks to predict SWE in the watersheds of Afghanistan. Marofi, Tabari, and Abyaneh (2011) use multivariate non-linear regression and

artificial neural networks with kriging. Meyal et al. (2020) present a recurrent neural network-based model using Long Short-Term Memory or LSTM (Hochreiter and Schmidhuber 1997) to predict SWE values in five SNOTEL locations in Colorado. In a concurrent work, Duan et al. (2023) (unpublished) compare multiple temporal deep learning models for SWE prediction, including LSTM, temporal convolutional neural network, and transformer architecture. In contrast to these efforts, *this* paper presents a modeling framework to capture attention in both spatial and temporal dimensions. Our evaluations show that attention can be effective, outperforming other deep learning approaches.

Approach

Let X denote the set of all n locations, and m be the number of days in any SWE season (e.g., ~ 270 days for Oct. 1 to July 1). Let S denote the set of all SWE seasons (i.e., years). Let $x_i^{j,h}$ denote a vector of directly observable attributes for location $i \in X$ on day $j \in [1, m]$ of season $h \in S$. Note that these attributes typically consist of static attributes (e.g., lat/long, elevation) and dynamic attributes (including daily weather variables, satellite reflectance observations). Finally, let $y_i^{j,h}$ be the corresponding daily SWE value, which is the primary prediction target in this work.

The goal of the prediction problem is: given the collection of attribute vectors $\{x_i^{j,h}\}$, enumerated over locations, seasons, and days, predict the corresponding daily SWE value of each $y_i^{j,h}$. An accurate solution to this problem allows for estimating any missing SWE value for a location, season, and day based on just observable attributes in the data without requiring any ground truth SWE. Training the model, however, does assume the availability of labeled SWE data.

Spatial Attention: We first consider our main ideas for implementing the spatial attention model. For mapping to the attention framework (Vaswani et al. 2017), we reformulate the prediction problem as a sequence-to-sequence prediction:

$$[y_1^{j,h}, y_2^{j,h}, \dots, y_n^{j,h}] \leftarrow f_{spatial}([x_1^{j,h}, x_2^{j,h}, \dots, x_n^{j,h}])$$

where, $f_{spatial}$ is the function for the transformer architecture to learn as described in more detail in Section Model Architecture.

There are several advantages to this formulation:

- First, the encoder within the transformer architecture can compute *attention for every pair of locations* (for the same day), providing an opportunity to exploit any spatial correlation between location pairs. Intuitively, this also implies that if there are missing data for some locations then the model may still be able to learn from other locations that are sufficiently correlated to the missing locations.
- Second, although the attention mechanism works in the spatial dimension, one can also *incorporate some temporal information*. More specifically, we divide each input location feature $x_i^{j,h}$ further into three parts i.e.; $x_i^{j,h} = [\phi, \alpha, \gamma]$. Here, ϕ represents the static spatial features

like elevation, lat/long, etc.; α represents the daily observations for the location (including weather and satellite image features); and γ represents the average of SWE observations on the same day j but from a window of adjacent w years, i.e. $\gamma = \frac{1}{2w+1} \sum_{t=h-w}^{h+w} x_i^{(j,t)}$. Intuitively, the idea of γ is to provide an added temporal context.

- There are other benefits to incorporating spatial attention. The set of *attention vectors* computed in the final model can be of value in itself, as it could reveal interesting hidden relations between the locations. Furthermore, in this formulation, each training example is for a given $\langle \text{day } j, \text{season } h \rangle$ combination—which implies at test time we can use this model as an *imputation tool* to fill SWE data for any location on arbitrary days of past SWE seasons.

Temporal Attention: The formulation for temporal attention in SWE prediction can be achieved using a simple modification to the spatial attention formulation. Intuitively, instead of posing each input as a sequence of locations, we pose it as a sequence of (m) days for each location i :

$$[y_i^{1,h}, y_i^{2,h}, \dots, y_i^{m,h}] \leftarrow f_{temporal}([x_i^{1,h}, x_i^{2,h}, \dots, x_i^{m,h}]),$$

where $f_{temporal}$ is the transformer-structured function to learn. Despite the similarity in the formulations, there are some subtle but important differences to the temporal attention scheme (compared to its spatial counterpart).

- First, unlike the spatial model (where we restrict α to a single day), here we have the opportunity to learn from a longer time window—e.g., a warmer-than-usual January and associated melt can change the snow density and age structure, which when followed by two cold months, can impact the SWE in a following warmer month. In the absence of any such distinctive long-term correlations, however, this architecture can be expected to perform similar to an RNN architecture such as LSTM.
- Furthermore, since the temporal attention scheme is tied to a single location (i) during prediction, any pattern to missing data (particularly from consecutive time windows) could potentially impact prediction accuracy.

Ensemble Attention: We can also combine the strengths of both attention mechanisms through ensemble methods. These methods can produce a net predicted value that could possibly surpass the prediction achieved individually. In principle, we can perform ensemble using any aggregator function \oplus :

$$y_{i,ensemble}^{j,h} = y_{i,spatial}^{j,h} \oplus y_{i,temporal}^{j,h}$$

We evaluate *average* as a simple aggregator, i.e., the average of the two predictions is the ensemble prediction.

Model Architecture

We adapt the transformer model architecture (Vaswani et al. 2017) to implement spatial and temporal attention. For ease of exposition, we describe the model architecture for the spatial scheme. Temporal attention can be implemented in the same architecture using a trivial modification to the inputs and outputs (as described above). Figure 3 shows our

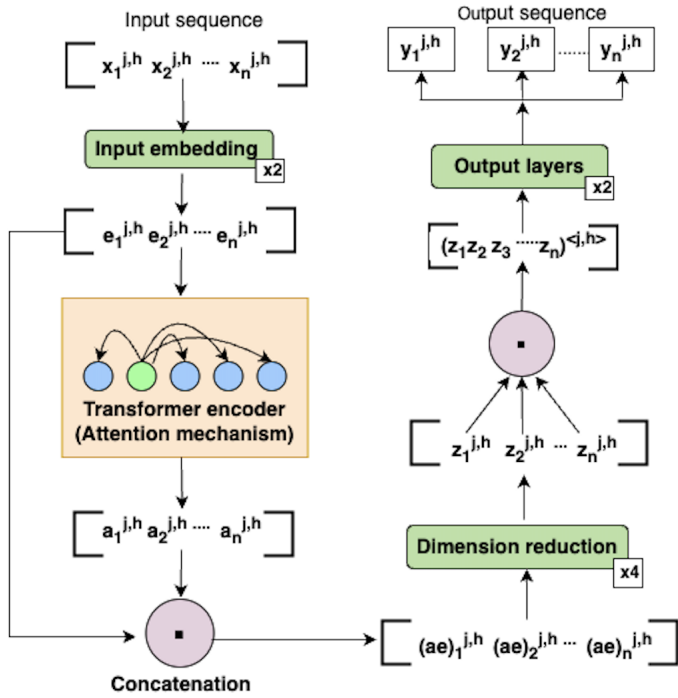


Figure 3: Spatial attention model architecture. This is easily adapted to also implement our temporal attention model.

transformer architecture for the spatial attention model. The model uses the classical transformer encoder (Vaswani et al. 2017) to encode attention between every pair of locations. The encoded information is then passed to a series of linear layers to obtain the SWE value for each location.

Input: Each training example is a sequence of locations $[x_1^{j,h}, x_2^{j,h}, \dots, x_n^{j,h}]$ for day j and season h . We denote this sequence as $\mathbf{x}^{j,h}$.

Output: The output of the model is the sequence of predicted daily SWE values $[y_1^{j,h}, y_2^{j,h}, \dots, y_n^{j,h}]$ for all locations for day j and season h . We denote this sequence as $\mathbf{y}^{j,h}$.

1. **Input Embedding Layers:** Two embedding linear layers are used to generate a d -dimensional embedding ($\mathbf{e}^{j,h}$) for the input sequence $\mathbf{x}^{j,h}$. We used $d = 512$ in our experiments. Both layers use Gaussian Error Linear Unit (GELU) for activation.
2. **Transformer Encoder:** For self-attention, we used 16 attention heads and 24 feed-forward layers within the block, to generate the encoded sequence ($\mathbf{a}^{j,h}$), which is $a_1^{j,h}, a_2^{j,h}, \dots, a_n^{j,h}$.
3. **Locationwise Concatenation:** Next, the encoded vector $\mathbf{a}^{j,h}$ is concatenated with the embedded vector $\mathbf{e}^{j,h}$, using a location-wise concatenation operation, represented by $\mathbf{ae}^{j,h}$: $(ae)_1^{j,h}, (ae)_2^{j,h}, \dots, (ae)_n^{j,h}$. The rationale behind this concatenation is to carry forward the initial input representation along with its encoded representation.
4. **Dimension Reduction Layers:** Note that the above step would have doubled the number of dimensions due to the

concatenation (i.e., $d = 1024$). Subsequently, each location’s representation in the sequence is passed through four linear layers to reduce its dimensionality by a factor of 8—i.e., if $d = 1024$ at the start of this step, it becomes 128 at the end. Furthermore, within each layer, a dropout with a rate of 0.20 is applied. The first and third layers were subjected to the Rectified Linear Unit (RELU) as an activation function. The dimension-reduced sequence is denoted as $\mathbf{z}^{j,h}$: $z_1^{j,h}, z_2^{j,h}, \dots, z_n^{j,h}$.

5. **All Locations Concatenation:** In this block, all elements in the sequence $\mathbf{z}^{j,h}$ are concatenated to form a single large vector of features, represented by $(z_1 z_2 z_3 \dots z_n)^{j,h}$. The rationale for this concatenation is that during the output extraction for an individual location, the output block will have an opportunity to make use of both its own location’s representation as well as others.
6. **Output Layers:** This block contains two linear layers that transform the concatenated representation to the final prediction for each of the n locations. The first layer uses GELU activation and has a dropout rate of 0.10. The second layer uses the identity activation function to obtain the daily SWE values ($\mathbf{y}^{j,h}$).
7. **Loss Function:** We used Mean Squared Error as the loss function, the AdamW Optimizer, and an initial learning rate of 0.0001. A scheduler reduces the learning rate by factor of 0.6 after every three epochs during training.

Implementation and Availability: Model implementation used the Python Scikit-learn (v1.1.2) (Linear Regression), and Pytorch (v2.0.1) (LSTM and Attention models) packages. Data processing and visualization used multiple Python and R packages. Code and data are accessible at: <https://github.com/Krishuthapa/SWE-Attention>.

Experimental Setup

Data Description: The set of static and dynamic features used along with their respective sources are listed below:

- **Static Features:** elevation, latitude, longitude (NRCS 2023); land cover (Yang et al. 2018); southness (NED 2014).
- **Dynamic (Daily) SNOTEL Features:** SWE, precipitation, minimum, maximum, and average temperatures (NRCS 2023)
- **Dynamic (Daily) Satellite Observations:** Passive microwave brightness temperature (19VE, 37VE, and their difference)(Brodzik et al. 2016)

The daily SNOTEL data consists of 822 stations for 18 water years (2002-2019).

We filtered out stations with more than 10% missing values in any variable or year. The resulting 323 stations comprised our main data set. Given our focus on SWE predictions, we used daily data for a 270 day-period starting Oct 1 as representative of the SWE season. This resulted in a total of 1,569,780 ($=323 \times 18 \times 270$) (location, year, day) combinations.

The satellite observations were available at spatial resolutions of 6.25 km (19 GHz) and 3.125 km (37 GHz). The 3.125 km resolution data was resampled to 6.25 km using a nearest neighbor interpolation. Values for a location

were assigned from the grid containing it. The landcover grid was at a 30m resolution and the dominant landcover type from the grid containing each location was extracted. A Digital Elevation Model 30m grid derived from the LIDAR point cloud source was processed in ArcGIS to obtain slope and elevation. From this southness was calculated as $\cos(\text{aspect}) \cdot \sin(\text{slope})$

Evaluation Methodology: We split the 18 years into two sets: training (13), and testing (5). The five testing years cover the entire spectrum of average annual SWE. Test water years are 2015, 2007, 2018, 2017, and 2008—ordered from the lowest to highest average annual SWE.

In our experiments, we compared five models:

- `Spatial_att`: the proposed spatial attention model
- `Temporal_att`: the proposed temporal attention model
- `Ensemble_att`: the proposed ensemble attention model which uses a simple average between the spatial and temporal attention predictions
- `LSTM`: Long Short-Term Memory model (Hochreiter and Schmidhuber 1997)
- `LR`: Linear regression

To compare models, we use the *Nash-Sutcliffe Efficiency* (*NSE*) metric (Nash and Sutcliffe 1970), a widely adopted standard for measuring the accuracy of hydrological models. *NSE* for each location $i \in X$ is calculated as follows:

$$NSE_i = 1 - \frac{\sum_{t=1}^T (Y_{a,i}^t - Y_{p,i}^t)^2}{\sum_{t=1}^T (Y_{a,i}^t - \bar{Y}_{a,i})^2} \quad (1)$$

where, T is the total number of distinct year-month-days used in testing for SWE prediction, $Y_{a,i}^t$ and $Y_{p,i}^t$ are actual and predicted SWE respectively for a given day $t \in [1, T]$, and $\bar{Y}_{a,i}$ is the mean of all the actual SWE values for location i over all T days. The value of *NSE* ranges from $-\infty$ to 1, with a value closer to 1 implying the best performance, and ≤ 0 implying that prediction is worse than using the long-term mean as the prediction. In practice, a value of 0.5 or above is subjectively considered good. We also calculated mean errors between the predicted and observed daily SWE values and annual maximum SWE values.

Experimental Results

Model Comparison: We first analyze the performance of models by their *NSE* values for the test dataset. *NSE* values were calculated for each location using Eq. 1. Figure 4 shows the distribution of all locations binned into different *NSE* groups. Observations are as follows.

- All three attention models outperform the other two models (`LSTM` and `LR`).
- Among the attention models, `Ensemble_att` is the best, followed closely by `Spatial_att` and then `Temporal_att`. For example, `Ensemble_att` predictions placed 82.04% of the locations with *NSE* values > 0.50 ; with the corresponding fractions being 65.63% for `Spatial_att` and 62.22% for `Temporal_att`.

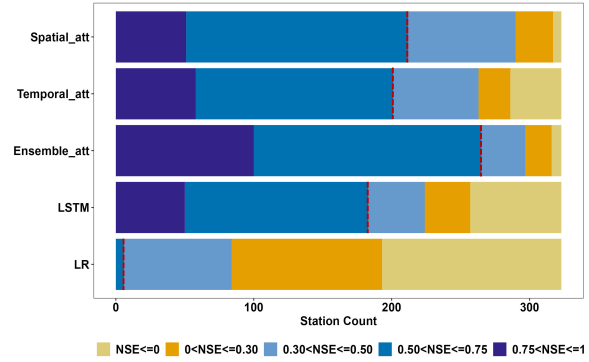


Figure 4: The distribution of all locations across five *NSE* groups for all models. Higher values (blue bars) indicate a better prediction.

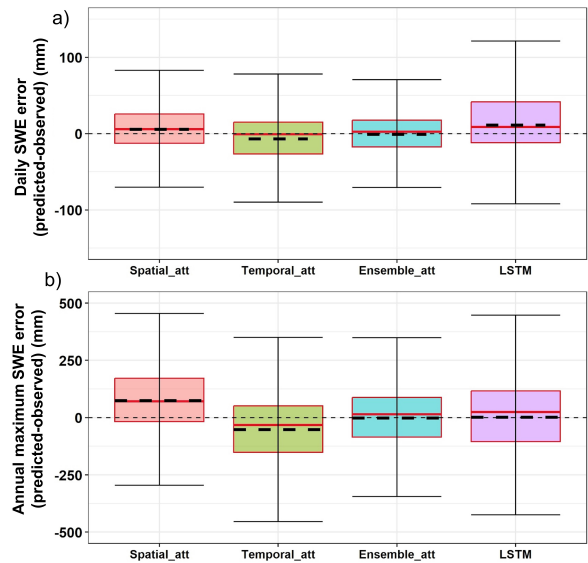


Figure 5: Mean error ($predicted - observed$) in SWE (mm) for each location. Part (a) errors in daily SWE across the 270 days of a SWE season. The range in each box plot corresponds to errors from 270 days \times 323 locations. Part (b) errors for the annual maximum SWE predictions.

- `Temporal_att` performs slightly better than `LSTM`. Even though both are temporal models, `LSTM` processes the sequence sequentially, while the attention mechanism can derive context from either end of the sequence.

Better performance of the `Ensemble_att` is partially due to the respective under- and over- prediction by the `Spatial_att` and `Temporal_att` models (Figure 5a), resulting in an ensemble with a higher accuracy and lower spread in errors. This behavior is pronounced when comparing the annual max SWE (Figure 5b).

While the `Ensemble_att` model works best in most locations, we observed a spatiotemporal variation in the best-performing model (Figure 6a). To examine model performance at a finer level, we plotted a relative performance

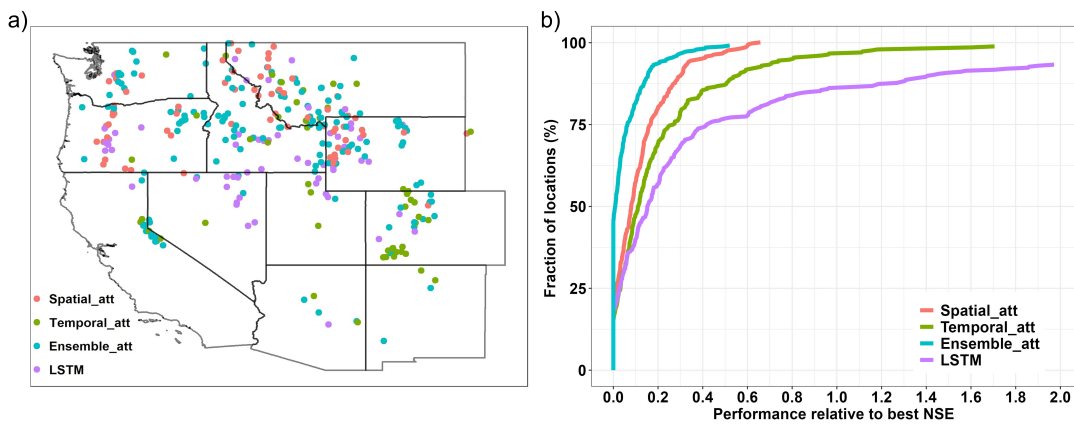


Figure 6: a) The best model (highest NSE) for each location overlaid on Western U.S. map. b) Relative model performance (RMP) chart: A model’s RMP is measured as the difference in its NSE from the best NSE for each location (X-axis). For visual clarity, the X-axis is restricted to a RMP of 2.0. This covers 94% of the data and primarily excludes LSTM data, resulting in the purple line not reaching a 100% fraction. The Y-axis is the fraction of locations for which that RMP is achieved. The closer a curve is to the Y-axis and for longer, the better.

chart shown in Figure 6b. For each location, we recorded the difference of the NSE obtained by a given model from the NSE for the best performing model for that location. The X-axis shows this difference as the “relative performance” (0 is best). We then plotted the fraction of locations (Y-axis) for which each model produced a given relative performance. In this representation, the closer a model’s curve aligns with the Y-axis the better. The best model, `Ensemble_att`, achieves the best NSE value for around 50% of the locations. Notably, even in cases where it is not the best, its difference from the best is at most 0.2 for almost 90% of the location. `Spatial_att` also provided a comparable performance.

Some of the spatial variation is attributable to differential performance for stations at different *elevations*—an important factor affecting snow dynamics. Table 1 shows the median NSE values within each elevation group for each model. In the lower elevation areas, which are generally harder to predict, `Spatial_att` performed significantly better. `Temporal_att` and `Ensemble_att` performed best at the highest and mid-range elevation groups respectively.

Model Run-Times: The models were trained on a 3.1GHz Apple M1 Pro processor with 16 GB RAM. The training times for the different models were: `Spatial_att` (440 min); `Temporal_att` (450 min); and LSTM (464 min). Note that the simple average ensemble can directly use the prediction outputs from both attention models.

Roadmap to Deployment

The objective of this work was to lay the groundwork for a generic and extensible attention-based ML modeling framework for SWE prediction. One of the main drivers behind this work is to develop capability to generate spatially-complete SWE maps at country-scales. Such maps can become an integral part of snow products useful to hydrologists, irrigation district managers, and other resource planning commissions and agencies such as the United States

Elevation Group	Spatial Att.	Temporal Att.	Ensemble Att.	LSTM
[0%, 25%)	0.50	0.24	0.49	0.31
[25%, 50%)	0.48	0.64	0.70	0.48
[50%, 75%)	0.61	0.66	0.72	0.59
[75%, 100%)	0.52	0.69	0.68	0.59

Table 1: Median NSE of each model across locations for four elevation groups. The groups are based on quartiles of elevations sorted from low to high elevation. Boldface entries have the highest median NSE for each group (row).

Bureau of Reclamation (USBR).

To enable this deployment, several extensions and evaluations will be needed including (but not limited to): a) application as a spatial imputation tool to fill SWE values for non-SNOTEL locations across past decades (so as to generate an extensive gridded product); and b) enable a forecast function (with uncertainty quantification) for prediction of future snowpack. Ideally such deployment should happen in close partnership with the public sector (e.g., a government agency or service center such as USDA NRCS or USBR).

Furthermore, this research has opened up multiple avenues of interesting research directions such as: i) non-trivial ensemble methods, ii) alternate ways to model the problem including spatiotemporal graph neural networks, or alternatively, use the attention scheme to infer a graph representation, iii) characterize the role of different variables, and iv) couple with process-based models to ensure scientific consistency and model interpretability.

Acknowledgments

This research was supported by USDA NIFA award No. 2021-67021-35344 (AgAID AI Institute). We thank Hossein Noorazar for help with some of the preliminary studies.

References

- Bair, E. H.; Abreu Calfa, A.; Rittger, K.; and Dozier, J. 2018. Using machine learning for real-time estimates of snow water equivalent in the watersheds of Afghanistan. *The Cryosphere*, 12(5): 1579–1594.
- Brodzik, M.; Long, D.; Hardman, M.; Paget, A.; and Armstrong, R. 2016. MEaSURES calibrated enhanced-resolution passive microwave daily EASE-grid 2.0 brightness temperature ESDR, version 1. *Digital Media*.
- Broxton, P.; Dawson, N.; and Zeng, X. 2016. Linking snowfall and snow accumulation to generate spatial maps of SWE and snow depth. *Earth and Space Science*, 3(6): 246–256.
- Carroll, S. S.; Carroll, T. R.; and Poston, R. W. 1999. Spatial modeling and prediction of snow-water equivalent using ground-based, airborne, and satellite snow data. *Journal of Geophysical Research: Atmospheres*, 104(D16): 19623–19629.
- Duan, S.; Ullrich, P.; Risser, M.; and Rhoades, A. 2023. Using Temporal Deep Learning Models to Estimate Daily Snow Water Equivalent over the Rocky Mountains. *Authorea Preprints*. Forthcoming.
- Elder, K.; Rosenthal, W.; and Davis, R. E. 1998. Estimating the spatial distribution of snow water equivalence in a montane watershed. *Hydrological Processes*, 12(10-11): 1793–1808.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Huang, H.-C.; and Cressie, N. 1996. Spatio-temporal prediction of snow water equivalent using the Kalman filter. *Computational Statistics & Data Analysis*, 22(2): 159–175.
- Kim, R. S.; Kumar, S.; Vuyovich, C.; Houser, P.; Lundquist, J.; Mudryk, L.; Durand, M.; Barros, A.; Kim, E.; Forman, B.; Gutmann, E.; Wrzesien, M.; Garneau, C.; Sandells, M.; Marshall, H.-P.; Cristea, N.; Pflug, J.; Johnston, J.; Cao, Y.; and Wang, S. 2021. Snow Ensemble Uncertainty Project (SEUP): Quantification of snow water equivalent uncertainty across North America via ensemble land surface modeling. *The Cryosphere*, 15: 771–791.
- Marofi, S.; Tabari, H.; and Abyaneh, H. Z. 2011. Predicting Spatial Distribution of Snow Water Equivalent Using Multivariate Non-linear Regression and Computational Intelligence Methods. *Water Resources Management*, 25: 1417–1435.
- Meyal, A. Y.; Versteeg, R.; Alper, E.; Johnson, D.; Rodzianko, A.; Franklin, M.; and Wainwright, H. 2020. Automated cloud based long short-term memory neural network based SWE prediction. *Frontiers in Water*, 2: 574917.
- Nash, J. E.; and Sutcliffe, J. V. 1970. River flow forecasting through conceptual models part I—A discussion of principles. *Journal of Hydrology*, 10(3): 282–290.
- NED, U. 2014. United States Geological Survey, National Elevation Dataset (<https://apps.nationalmap.gov/datasets/>). Accessed: November 3, 2023.
- NRCS. 2023. USDA NRCS NWCC Database for SNOTEL locations data (<https://wcc.sc.egov.usda.gov/reportGenerator/>). Accessed: November 3, 2023.
- Rutter, N.; Essery, R.; Pomeroy, J.; and Altimir, N. 2009. Evaluation of forest snow processes models (SnowMIP2). *Journal of Geophysical Research: Atmospheres*, 114(D6).
- Sospedra-Alfonso, R.; Melton, J.; and Merryfield, W. 2015. Effects of temperature and precipitation on snowpack variability in the Central Rocky Mountains as a function of elevation. *Geophysical Research Letters*, 42: 4429–4438.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Yang, L.; Jin, S.; Danielson, P.; Homer, C.; Gass, L.; Bender, S. M.; Case, A.; Costello, C.; Dewitz, J.; Fry, J.; Funk, M.; Granneman, B.; Liknes, G. C.; Rigge, M.; and Xian, G. 2018. A new generation of the United States National Land Cover Database: Requirements, research priorities, design, and implementation strategies. *ISPRS Journal of Photogrammetry and Remote Sensing*, 146: 108–123.