

# Multi-step Weekly Average Forecasting of Reservoir Storage Volume Using Deep Learning

Zachary Herbert<sup>1</sup>, Zeeshan Asghar<sup>2</sup>, Carlos A. Oroza<sup>2</sup>

<sup>1</sup>Central Utah Water Conservancy District

<sup>2</sup>Dept. Civil and Environmental Engineering, University of Utah

## Key Points:

- A deep learning algorithm is developed for multi-step reservoir storage volume forecasting from snow water equivalent.
- The algorithm has higher forecast accuracy compared to common statistical methods (SARIMA, VAR, and TBATS).
- The algorithm performed best for years with large runoff; worst for years with small runoff and late-season snow accumulation.

---

Corresponding author: Zachary Herbert, [zachary@cuwcd.com](mailto:zachary@cuwcd.com)

13 **Abstract**

14 Machine-learning algorithms have shown promise for streamflow forecasts, reser-  
15 voir operations, and scheduling, but have exhibited lower accuracy in predicting extended  
16 time horizons of peak storage volume (PSV). Deep learning algorithms exhibited improved  
17 inflow forecasting accuracy, but existing research has been mostly limited to real-time  
18 operation and short-term planning. We evaluate a new approach based on a hybrid ResCNN-  
19 LSTM Encoder-Decoder algorithm, enabling long-term multi-step reservoir forecasts. The  
20 proposed approach provides a three-month, weekly averaged prediction of reservoir stor-  
21 age volume (RSV) during the runoff season based on historical snow water equivalent  
22 (SWE). The optimal architecture and hyper-parameters for the model are configured through  
23 five-fold cross validation resulting in a twelve-layered residual convolutional neural net-  
24 work (ResCNN) as the encoder and a four-layered long short-term memory (LSTM) neu-  
25 ral network as the decoder. We evaluate the algorithm using 30 years of RSV and SWE  
26 data at the Upper Stillwater Reservoir located in Utah. The most accurate long-term  
27 predictions occurred during periods of large runoff (in excess of 28,000 ac-ft). The pe-  
28 riods where the model performed the worst were during small runoff and late-season SWE  
29 accumulation. We find that the ResCNN-LSTM consistently outperforms three widely  
30 used statistical models, with an average PSV absolute percent error of 2.66% for the pro-  
31 posed algorithm compared to SARIMA (14.22%), TBATS (13.82%), and VAR (18.14%).

## 1 Introduction

Classical statistical models such as Seasonal Auto-regressive Integrated Moving Average (SARIMA) (Papamichail & Georgiou, 2001), Vector Auto-regression (VAR) (Iddrisu et al., 2016) and Trigonometric Seasonal Box-Cox Transformation with ARMA residuals, Trend, and Seasonal Components (TBATS) (Elizaga et al., 2014) have long been employed for reservoir storage and outflow prediction. These models are well-suited to short-term forecasts, but have limited capacity for long-term forecasts due to the convergence of the auto-regressive part of the model to the mean of the time series (Shumway & Stoffer, 2000). Forecast skill is also confounded by hydro-meteorological predictability in snow-dominated catchments (Anghileri et al., 2016). Machine-learning algorithms provide an alternative approach and are increasingly being used in a variety of related hydrologic fields: rainfall-runoff prediction for ungauged basins (Kratzert et al., 2019), hydropower production forecasting (Stokelj et al., 2002), spatial snow water equivalent (SWE) estimation for mountainous areas (Zheng et al., 2018), and quantifying climate and catchment control on hydrological drought (Konapala & Mishra, 2020).

Machine learning has seen broad application to reservoir forecasting in both direct and multi-step scenarios, and generally resulted in more reliable forecasts of inflow extremes. Coulibaly et al. (2000) trained a feed-forward neural network using an early stopping training approach for real-time reservoir inflow forecasting with lead times of one to seven days. An improvement to daily reservoir inflow forecasts was later made using a robust weighted-average ensemble that takes advantage of three different models: nearest neighbors, a conceptual model, and an artificial neural network (Coulibaly et al., 2005). An additive ensemble for monthly reservoir inflow forecasting was developed by Bai et al. (2015), incorporating an auto-regressive model, least squares support vector machine, and adaptive neuro-fuzzy inference system to subforecast trend, period, and stochastic terms. Bourdin et al. (2014), Wang et al. (2012), and Ahmed et al. (2015) all used ensemble learning methods coupled with meteorological predictions to forecast reservoir inflows with respective forecast horizons of three, eight, and fourteen days. Long-range streamflow forecasts, extending twelve months, were assessed by Bennett et al. (2016) in which calibrated climate forecasts are combined with a conceptual runoff model and a three-staged error model to simulate reservoir inflows. Similarly, Y. Liu et al. (2017) developed a long-term streamflow forecasting scheme, extending nine months, utilizing

64 Random Forest and support vector regression for precipitation post-processing of numer-  
65 ical weather predictions to feed into a hydrological watershed model.

66 Research with deep learning algorithms applied to reservoir inflow forecasting has  
67 found improved forecast accuracy, but has been mostly limited to real-time operation  
68 and short-term planning. Deep belief networks for multiscale feature learning (Bai et al.,  
69 2016) improved on prior efforts (Bai et al., 2015) in direct-step forecasting using an ad-  
70 ditive ensemble approach. Budu (2014) and Chiamsathit et al. (2016) both applied multi-  
71 layered perceptrons for direct-step forecasting scenarios that achieved reasonable accu-  
72 racy among daily and monthly timesteps, respectively. Time-lagged recurrent neural net-  
73 works (TLRNs) have been studied (Muluye & Coulibaly, 2007; Kote & Jothiprakash, 2009;  
74 Sattari et al., 2012) where a preceding record of reservoir inflow is used to investigate  
75 the performance of a back-propagation through time (BPTT) algorithm. Better predic-  
76 tion of inflow into a reservoir using TLRN was achieved by Kote and Jothiprakash (2009)  
77 by modifying the artificial neural network to include seasonal (monsoon) effects, accu-  
78 rate mapping of high and low flows was achieved following a monthly time-step.

79 However, attempts at long-range forecasts with deep learning algorithms have gener-  
80 ally exhibited lower accuracy. Similar to Bennett et al. (2016) and Y. Liu et al. (2017),  
81 earlier attempts at long-term forecasting (Muluye & Coulibaly, 2007; Kote & Jothiprakash,  
82 2009) failed to accurately predict the peak storage volume (PSV) at extended time hori-  
83 zons. More recently, multi-step flood forecast models (Chang & Tsai, 2016; Zhou et al.,  
84 2019; Kao et al., 2020) have been developed for predicting reservoir inflows using adap-  
85 tive neuro-fuzzy inference systems and long short-term memory (LSTM) based Encoder-  
86 Decoder frameworks. However, the timestep for these forecasts are hourly with forecast  
87 horizons only extending four, six, and eight hours.

88 We evaluate a new approach to multi-step weekly average forecasting of reservoir  
89 storage volumes (RSV) based on historical regional SWE data. This approach aims to  
90 use existing SWE and RSV time series data to train a hybrid, multi-step ResCNN-LSTM  
91 (Encoder-Decoder) (LeCun, 1990; Hochreiter & Schmidhuber, 1997) algorithm, to pre-  
92 dict the RSV of future timesteps for the proceeding three months. In contrast to RNNs,  
93 convolutional neural networks (CNN) operate independently of previous time steps to  
94 capture fixed size contexts, allowing for parallel computation within a given sequence.  
95 The stacking of convolutional layers allows for precise control of the dependencies to be

modeled by effectively increasing the context size (Gehring et al., 2017). Implementing residual connections in CNNs has been shown to improve model performance by increasing the depth of the model architecture (He et al., 2015). Deep learning frameworks utilizing residual connections in CNNs have seen extensive application in other fields (H. Liu & Song, 2018; Ning et al., 2019; Cengil & Cinar, 2018), but have not yet been evaluated in the context of RSV forecasting. Thus, a key distinction between the proposed model in this paper and the others listed above is that a deep residual CNN (ResCNN) is initially used for feature extraction, which may improve long-range forecast accuracy. The specific aims of this study are:

1. Assess the accuracy of the ResCNN-LSTM algorithm in predicting RSV from SWE.
2. Determine characteristics of years in which the algorithm exhibits high vs. low accuracy.
3. Compare the long-term accuracy of the proposed algorithm against three commonly used RSV forecast methods: SARIMA, VAR, and TBATS.

## 2 Materials and Methods

### 2.1 Study Region

This study focuses on the Upper Stillwater reservoir located at the top of the Central Utah Water Conservancy District’s (CUWCD) collection system in the Uinta Mountains. CUWCD is one of Utah’s four large specialty water districts that provides potable and secondary water to various water associations, conservancy districts, irrigation companies, and local residents. The water district spans eight counties with over \$3.5 billion in infrastructure. There are currently ten lakes/reservoirs maintained and operated by CUWCD that house non-potable water in excess of 1.6 million ac-ft. The storage levels for these reservoirs act as a barometer for the state’s water resources and provide insight for how to appropriately prepare for future water usage. Figure 1 shows Upper Stillwater (located in the middle of the figure) surrounded by a network of snow telemetry monitoring sites.

### 2.2 Data

The data used for model training is accessed from the National Resource Conservation Service (NRCS) online Application Programming Interface (NRCS, 2020 (accessed

126 May 3, 2020)). The two data types used for this study are RSV and SWE (the depth  
127 of water (in) that would result if the snowpack were melted). Specific to Upper Stillwa-  
128 ter Reservoir, the historical daily data span from January 1990 to the present with new  
129 values updated daily. Linear time interpolation was used for gap-filling a limited num-  
130 ber of missing data points within the time series. From January 1st 1990 to July 14th  
131 2019, a total of 16 missing daily data points, out of 10,787, required interpolation. These  
132 points occurred in 1990, during the first year of data collection.

133 The RSV time series for Upper Stillwater illustrates a seasonal runoff period, fed  
134 by snowmelt, that begins in April and ends in July. A governing assumption of this study  
135 is that water managers only have until the end of March to make a final decision regard-  
136 ing the level of storage space to leave vacant in the reservoir for the runoff season. There-  
137 fore, the critical period to forecast RSV is a 15-week window between the first of April  
138 and the beginning of July. Such an extended forecast horizon requires a model capable  
139 of learning long-term dependencies.

140 SWE data were collected from the NRCS monitoring network of snow telemetry  
141 sites for the same period as Upper Stillwater’s RSV time series. For each of the avail-  
142 able monitoring sites, a maximum of three daily data points required interpolation over  
143 the entire period. The dependence of RSV (Figure 2) on SWE (Figure 3) is the primary  
144 relationship that the model will attempt to learn.

145 The daily data were then prepared for training the algorithm. The data were first  
146 re-sampled into weekly averages and scaled between 0 and 1 based on the chosen acti-  
147 vation function for the model (see Section 2.3). A variety of sliding window lengths (15,  
148 20, and 25 weeks) were then used as inputs to predict the next 15 weeks. This range of  
149 window sizes is selected based on the time series data for SWE (depending on the pre-  
150 cipitation distribution during a given winter season, the process of SWE accumulation  
151 ranges between 15 to 25 weeks). The input window length that yields the greatest per-  
152 formance is selected for the final model.

### 153 **2.3 Deep Learning Model**

154 The objective of the model is to forecast multiple timesteps forward based on mul-  
155 tiple inputs from the past. The inputs are the multiple time series of RSV and SWE and  
156 the output is a future RSV sequence prediction starting at the final point in the input

157 data. Therefore, a multivariate sequence-to-sequence prediction model is required. This  
158 type of model is broken down into two separate models: one for reading the input se-  
159 quence and encoding it into a fixed-length vector (Encoder), and a second model for de-  
160 coding the fixed-length vector and outputting the predicted sequence (Decoder) (Sutskever  
161 et al., 2014). Following the Decoder, a time-distributed fully connected layer is used as  
162 the final component to condense the output and yield a forecasted sequence of values.

163 Figure 4 illustrates the architecture of this model, which receives a sliding window  
164 input of multiple variables and transmits a sliding window output of a single target vec-  
165 tor. This architecture has proven to be effective for numerous sequence-to-sequence prob-  
166 lems, including multi-step flood forecasting (Kao et al., 2020), network traffic forecast-  
167 ing (Zhang & You, 2020), weather forecasting (Yuan et al., 2019), and predicting solar  
168 performance ratio (Yen et al., 2019). The Encoder-Decoder model is written in Python  
169 using the Keras deep-learning library (Gulli & Pal, 2017). Another important feature  
170 of the proposed model is how each node sequentially transmits information to the next  
171 within each layer. This process is governed by the use of a piece-wise linear activation  
172 function. A node or unit that implements this activation function is referred to as a rec-  
173 tified linear unit (ReLU) (Agarap, 2018). The benefit of using ReLu is two-fold: it al-  
174 lows for faster training time due to its near-linear properties, while also addressing the  
175 vanishing gradient problem during back-propagation of errors. The Adam optimizer (Kingma  
176 & Ba, 2014) is used to adaptively optimize the weights within the network using con-  
177 cepts of momentum (Sutskever et al., 2013) and stochastic gradient descent (Robbins  
178 & Monro, 1951).

### 179 **2.3.1 CNN Encoder**

180 A deep ResCNN is used as the encoder in the Encoder-Decoder architecture. A one-  
181 dimensional CNN is a model with one or more hidden layers that operate over a 1D se-  
182 quence (e.g. sentence or time series) through convolutions. In a multi-layer CNN, the  
183 stacking of CNN layers creates a hierarchical structure that provides a shorter path to  
184 capture long-range dependencies. Therefore, the model creates hierarchical representa-  
185 tions over the input sequence allowing nearby input elements to interact at lower lay-  
186 ers while distant elements interact at higher layers (Gehring et al., 2017). Causal padding  
187 (van den Oord et al., 2016) is used for each CNN layer to ensure the model does not vi-  
188 olate the temporal order (i.e. model does not have look-ahead bias). Residual connec-

189 tions are used between stacked CNN layers to boost model performance with increased  
190 network depth (He et al., 2015). A residual block is defined with a single skip connec-  
191 tion between CNN layers (Figure 5). Within a residual block, the output from the first  
192 CNN layer is passed through the activation function (ReLU) and sent in two separate  
193 paths: forward to the next CNN layer in sequence, and around to skip the next CNN  
194 layer as a residual from the previous. The residual is added to the output of the second  
195 CNN layer prior to its respective output being activated through ReLU. Thus, the model  
196 is optimized for a residual mapping of feature extraction from the sliding window of time  
197 series input.

198 The encoder is designed to gradually reduce the dimensionality of the input fea-  
199 ture matrix while increasing the number of feature abstractions. This is done using fil-  
200 ters (see Section 2.3.3) and pooling layers, whose purpose is to condense a CNN layer’s  
201 output to the most prominent elements. Max pooling is used at the end of each resid-  
202 ual section (Figure 5), in which two residual blocks are connected in sequence for a given  
203 number of filters and kernel size. Max pooling and flatten are used at the end of the en-  
204 coder to downsize the extracted features into a fixed length vector proportional to the  
205 number of nodes in the decoder. The final output represents the extracted elements as  
206 features from the input sequences that will be fed as a flattened sequence for the decoder.

### 207 **2.3.2 RNN Decoder**

208 The extracted features from the encoder are fed into the decoder to yield a fore-  
209 casted sequence of values. This is done using an RNN capable of learning long-term de-  
210 pendencies, which is analogous to context for the case of sequence-to-sequence predic-  
211 tion. RNNs are well suited for time series data as they process each timestep sequen-  
212 tially for modelling non-linear relationships between the input and the output. This is  
213 achieved by forming recurrent cycles within the nodes/cells of each hidden layer (Kao  
214 et al., 2020). In the case of runoff from snowmelt, a typical approach for determining the  
215 type of runoff season to expect is to compare the current snowpack with previous years.  
216 The process of observing current data and recalling previous significant events and their  
217 outcomes is what the RNN attempts to mimic. However, RNNs fail to connect informa-  
218 tion from the past (input) to the present (output) data when the gap between the two  
219 grows too large, which creates the issue of long-term dependencies. This problem is ad-  
220 dressed with the use of LSTM networks (Hochreiter & Schmidhuber, 1997), a special kind

221 of RNN. An essential feature of the LSTM cell is its state that runs directly through the  
222 network, enabling addition or removal of information from the cell state via regulated  
223 gates. These internal gates are weighted functions that govern how the information flows  
224 in the cell, mitigating the vanishing gradient problem.

### 225 *2.3.3 Model training and assessment*

226 The optimal hyper-parameters for the model (i.e. layers, nodes, filters, kernels, in-  
227 put window, epochs, and batch size, see Table 1) are determined through five-fold cross  
228 validation. Layers are arrays of nodes that sequentially transmit information from one  
229 to the next. Within each layer are nodes connected by multiple weights for a given num-  
230 ber of inputs and outputs. A single node receives input data, processes the input as a  
231 weighted sum, and then propagates new information to its successor based on a given  
232 activation function. Filters and kernels are interrelated hyper-parameters specific to CNNs.  
233 A kernel represents a matrix of weights that slide over the input sequence calculating  
234 the dot-product between the sequence values and matrix weights. Therefore, the size of  
235 the kernel represents the length of the window it spans for deep feature extraction. A  
236 complete tour of a kernel over the input sequence represents a filter; thus, kernel's op-  
237 erating over multiple channels of input establish a filter/feature map. The input win-  
238 dow represents the multivariate time series of RSV and SWE accumulation during the  
239 winter season until the first of April. Epochs represent the number of full passes of the  
240 data set that the model uses during training. The batch size is the fraction of data that  
241 the model is exposed to during each epoch. An early stopping algorithm is used during  
242 training to prevent over-fitting with excessive epochs: a training session will terminate  
243 early if there are 10 consecutive epochs with no improvement in minimizing the mean  
244 squared error.

245 For a given input window length, the training data for each test year spans from  
246 Jan 1990 to mid October or November of the previous year. For example, the training  
247 data set for the 2015 runoff period with a 20-week input window begins January 1990  
248 and ends November 2014. 80% of the data are used for training, while 20% are used for  
249 validation. This inner split between training and validation data allows for learning curves  
250 to be developed (to evaluate signs of over vs under-fitting). Across all five test years, the  
251 model is trained one configuration at a time to develop the regression metrics: mean ab-  
252 solute error (MAE), root mean squared error (RMSE), median absolute error (MedAE),

253 Nash-Sutcliffe model efficiency coefficient (NSE), and explained variance (ExpVar). The  
254 regression metrics are calculated from a held-out data set that is not used during model  
255 training. The hold-out set spans 30-40 weeks and consists of two consecutive parts: in-  
256 put (15-25 weeks) and output (15 weeks). The input spans from October or November  
257 through May; the output from April through June. For example, the hold-out data set  
258 used to forecast the 2015 runoff period with a 20-week input window begins in Novem-  
259 ber 2014 and ends at the start of April 2015.

260 Finally, a confidence interval is calculated from multiple model runs. Due to the  
261 stochastic nature of the model, a slightly different forecast will be returned each time  
262 the model is trained. Therefore, for a given configuration, the model is trained multi-  
263 ple times to establish a normal distribution of model predictions for each time-step in  
264 the forecast. By design, forecast points will be labeled as outliers if they lie beyond the  
265 whiskers of their respective boxplot following the Tukey method (Tukey, 1970).

266 The forecast plots and regression metrics provide valuable insight into the overall  
267 accuracy of the ResCNN-LSTM; however, the primary concern is the total expected runoff  
268 (TER). This can be defined as the change in RSV from the end of March to the begin-  
269 ning of July. The TER is the amount of water that is expected to fill the reservoir dur-  
270 ing the critical runoff period. Hence, the statistics of greatest value lie at the crest of each  
271 forecast curve (i.e. the PSV).

## 272 **2.4 Statistical Methods for Comparison**

273 The forecasts are compared against three widely used statistical models: Seasonal  
274 Auto-regressive Integrated Moving Average (SARIMA), Vector Auto-regression (VAR),  
275 and Trigonometric Seasonal Box-Cox Transformation with ARMA residuals, Trend, and  
276 Seasonal Components (TBATS). Each model is trained on monthly averaged data; there-  
277 fore, their forecasts for PSV are the maximum monthly average during the runoff pe-  
278 riod. This change in the timestep frequency is due to the limitations of the statistical  
279 models to forecast into such extended horizons.

280

### 2.4.1 SARIMA Model

281

282

A discrete time series  $Z_1, Z_2, Z_3, \dots, Z_{N-1}, Z_N$  of measurements at equal time intervals is simulated by a stochastic SARIMA model (Box et al., 2015) given by:

$$\varphi(B)\Phi(B^S)(1-B)^d(1-B^S)^{DZ_t} = \theta(B)\Theta(B^S)e_t \quad (1)$$

283

284

285

286

287

288

Here,  $t$  represents the discrete time and  $S$  denotes the length of each season. The  $B$  term corresponds to the backward shift operator which is defined by  $BZ_t = Z_{t-1}$  and  $B^SZ_t = Z_{t-S}$ . The independently and normally distributed white noise residual is represented by  $e_t = [\text{NID}(0, \sigma_e^2)]$  which has a zero mean, and variance defined by  $\sigma_e^2$ . From the left hand side of equation 1, the first two terms  $\varphi$  and  $\Phi$  represent series expansions given by:

289

$$\varphi(B) = 1 - \varphi_1B - \varphi_2B^2 - \dots - \varphi_pB^p \quad (2)$$

$$\Phi(B^S) = 1 - \Phi_1B^S - \Phi_2B^{2S} - \dots - \Phi_pB^{pS} \quad (3)$$

290

291

292

293

294

295

296

297

298

Equation 2 represents the nonseasonal auto-regressive operator of order  $p$  and  $\varphi_i, i = 1, 2, \dots, p$  depicts the nonseasonal auto-regressive parameters.  $(1-B)^d$  is the nonseasonal difference operator of order  $d$  which produces nonseasonal stationarity of the  $d$ th differenced data, usually  $d = 0, 1$ , or  $2$ . Whereas, equation 3 depicts the seasonal auto-regressive operator of order  $P$  and  $\Phi_i$ ; herein,  $i = 1, 2, \dots, P$  are the seasonal auto-regressive parameters.  $(1-B^S)^D$  is the seasonal differencing operator of the order  $D$  to produce seasonal stationarity of the  $D$ th differenced data, usually in the order of  $D = 0, 1$ , or  $2$ . From the right hand side of equation 1, the first two terms  $\theta$  and  $\Theta$  represent series expansions given by:

299

$$\theta(B) = 1 - \theta_1B - \theta_2B^2 - \dots - \theta_qB^q \quad (4)$$

$$\Theta(B^S) = 1 - \Theta_1B^S - \Theta_2B^{2S} - \dots - \Theta_QB^{QS} \quad (5)$$

300

301

302

303

Equation 4 is the nonseasonal moving average operator of the order  $q$ ; thus, equation 4 and  $q$  are the nonseasonal moving average parameters. Equation 5 is the seasonal moving average operator of order  $Q$  and  $\Theta_i, i = 1, 2, \dots, Q$  are the seasonal moving average parameters. Lastly, the natural log of the RSV time series is taken to stabilize the

304 variance of the time series and to transform any skew in the distribution into a normal  
 305 distribution (Papamichail & Georgiou, 2001). Using an annual seasonal term (S) of 12  
 306 months, the SARIMA model parameters (p,d,q)(P,D,Q) are configured using the autoarima  
 307 function provided in the Pmdarima statistical library in Python 3.0.

### 308 **2.4.2 VAR Model**

309 VAR is another frequently used model for multivariate time series. The basic VAR  
 310 model of order p, as suggested by Sims (1980) is given by

$$y_t = A_1 y_{t-1} + A_2 y_{t-2} + \dots + A_p y_{t-p} + CD_t + u_t \quad (6)$$

311 Where  $y_t = (y_{1t}, y_{2t}, \dots, y_{Kt})'$  represents a vector of K observable endogenous variables  
 312 and  $D_t$  consists of all deterministic variables which carry a constant, a linear trend, sea-  
 313 sonal dummy variables and user-specified variables.  $u_t$  is a K-dimensional unobservable,  
 314 zero-mean, white noise process which has a positive definite co-variance matrix  $E(u_t u_t') =$   
 315  $\sum_u \cdot A_i$  and C are parameter matrices of suitable dimension upon which various restric-  
 316 tions can be imposed. For a K-dimensional auto-regression with an effective sample size  
 317 N, the optimal lag order p is selected that minimizes the Akaike Information Criteria (AIC)  
 318 given by the following equation:

$$AIC(p) = \ln|\bar{\sum}(p)| + \frac{2}{N}(K^2 p) \quad (7)$$

319 The  $\bar{\sum}(p)$  is the quasi-maximum likelihood estimate of the innovation covariance ma-  
 320 trix  $\sum(p)$  (Ivanov & Kilian, 2005; Sin & White, 1996). The parameters in equation 6  
 321 are estimated by the method of generalized least squares. This is done by first estimat-  
 322 ing the individual equations of the system by ordinary least squares. The residuals can  
 323 then be utilized to estimate the white noise co-variance matrix  $\sum_u$  as  $\widehat{\sum}_u = T^{-1} \sum_{t=1}^T \widehat{u}_t \widehat{u}_t'$   
 324 which is used to compute the generalized least square estimator (Iddrisu et al., 2016).  
 325 The VAR model is developed through the statsmodels Python module and utilizes the  
 326 same SWE data as the ResCNN-LSTM.

### 327 **2.4.3 TBATS Model**

328 BATS is a combination of three methodologies; (i) Exponential Smoothing Method,  
 329 (ii) Box-Cox Transformation and (iii) ARMA model for residuals. The Box-Cox Trans-  
 330 formation helps to deal with non-linear data and ARMA model for residuals can de-correlate

331 the time series data. However, the BATS model does not do well when the seasonality  
 332 is complex with high frequency. Thus, De Livera et al. (2011) proposed a TBATS model,  
 333 which includes the trigonometric seasonal component. The trigonometric expression of  
 334 seasonality terms serves to reduce the parameters of model when seasonal frequencies  
 335 (e.g. annual streamflow or annual peak flow) are high and improves the model flexibil-  
 336 ity (i.e. lower bias with higher variance), enabling it to handle complex seasonality.

337 The TBATS model is comprised of the following terms:

$$\begin{aligned}
 y_t^{(\lambda)} &= l_{t-1} + \phi b_{t-1} + \sum_{i=1}^T s_{t-m_i}^{(i)} + d_t \\
 l_t &= l_{t-1} + \phi b_{t-1} + \alpha d_t \\
 b_t &= \phi b_{t-1} + \beta d_t \\
 d_t &= \sum_{i=1}^p \varphi_i d_{t-i} + \sum_{i=1}^q \theta_i e_{t-i} + e_t
 \end{aligned}$$

338 Where:

339  $y_t^{(\lambda)}$  – time series at moment  $t$  (Box-Cox transformed)

340  $s_t^{(i)}$  –  $i$  th seasonal component

341  $l_t$  - local level

342  $b_t$  - trend with damping

343  $d_t$  –  $ARMA(p, q)$  process for residuals

344  $e_t$  - Gaussian white noise

345 Seasonal part:

$$\begin{aligned}
 s_t^{(i)} &= \sum_{j=1}^{(k_i)} s_{j,t}^{(i)} \\
 s_{j,t}^{(i)} &= s_{j,t-1}^{(i)} \cos(\omega_i) + s_{j,t-1}^{*(i)} \sin(\omega_i) + \gamma_1^{(i)} d_t \\
 s_{j,t}^{*(i)} &= -s_{j,t-1}^{(i)} \sin(\omega_i) + s_{j,t-1}^{*(i)} \cos(\omega_i) + \gamma_2^{(i)} d_t \\
 \omega_i &= 2\pi j/m_i
 \end{aligned}$$

346 Model parameters:

347  $T$  - Amount of seasonalities

348  $m_i$  - Length of  $i$ th seasonal period

349  $k_i$  - Amount of harmonics for  $i$ th seasonal period

350  $\lambda$  - Box-Cox transformation

351  $\alpha, \beta$  – Smoothing

352  $\phi$  - Trend damping

353  $\varphi_i, \theta_i$  –  $ARMA(p, q)$  coefficients

354  $\gamma_1^{(i)}, \gamma_2^{(i)}$  - Seasonal smoothing (two for each period)

355 Based on a Fourier series, each seasonality is modeled by a trigonometric represen-  
356 tation. Advantages of this model framework consist of (i) improved parameter estima-  
357 tion and (ii) an ability to handle complex components and features intrinsic to the time  
358 series. As an innovations state-space model, TBATS admits a larger parameter space  
359 with the possibility of better forecasts (Hyndman, 2008). Likewise, TBATS involves a  
360 much simpler and efficient estimation procedure than other state-space models (e.g. SARIMA  
361 and VAR). With the use of trigonometric functions, the model is capable of handling non-  
362 integer seasonal frequencies (e.g. 365.25 days in a year), while also accommodating both  
363 nested and non-nested seasonal components. Additionally, the model handles nonlinear  
364 features typically seen in time series while taking into account any auto-correlation within  
365 the residuals (De Livera et al., 2011). The TBATS model is implemented using the `tbats`  
366 package in Python 3.0 and incorporates quarterly, bi-annual, and annual seasonal pe-  
367 riods.

### 368 **3 Results**

#### 369 **3.1 Model Architecture**

370 The optimal architecture for the ResCNN-LSTM was found to be a twelve-layered  
371 ResCNN as the encoder and a four-layered LSTM as the decoder. Through five-fold cross  
372 validation, the optimal number of filters in each ResCNN section was found to be 16 (fil-  
373 ter1), 32 (filter2), and 64 (filter3), while the LSTM layers required 64 nodes/cells per  
374 layer. The time-distributed fully connected layer is configured with 32 nodes to begin  
375 condensing the output from the decoder. Table 1 outlines the optimal configuration and  
376 architecture for the model. The maximum number of epochs is set to 100, but the av-  
377 erage final converged number ranges between 90 and 100 for a given test year.

#### 378 **3.2 Model Performance**

379 This section is organized by each test year starting from 2015 and ending in 2019.  
380 During each test year, the date range for each time series input stems from the second  
381 week in November through the end of March. Figure 6 illustrates the forecasted runoff  
382 period which is composed with the actual values (solid green), forecasted values (dot-  
383 ted red), reservoir spill limit (dashed black), 95% confidence interval (shaded grey) and  
384 boxplots. The box plots and regression metrics are based on a sample population of 30

385 individual model trainings per test year. Regression metrics for the ResCNN-LSTM fore-  
386 casts are shown in Table 2 below. Results from the analysis of statistical models are shown  
387 in Figure 7. The graphs represent the storage volume forecasts for all three models across  
388 the five consecutive test years.

389       Regardless of the low snowpack, the 2015 forecast proved to be very accurate. In  
390 addition to the low errors summarized in Table 2, the 15-week forecast nearly managed  
391 to identify the exact week that the reservoir was going to fill. The forecast’s accuracy  
392 is highest near the end of the period and lowest in the beginning. The 2016 forecast is  
393 an improvement over the previous year, with smaller errors for the MAE, RMSE, and  
394 MedAE, and improved values for NSE and ExpVar. The 2015 and 2016 forecasts are more  
395 skillful in the long-term projections occurring after May rather than earlier in April. Among  
396 all of the test years, the forecasts for 2015, 2016 and 2019 were the most accurate with  
397 the lowest reported errors. The forecast for 2017 fails to identify when the RSV will first  
398 increase, along with when the storage volume will peak. However, the forecast still man-  
399 ages to achieve an accurate prediction for the final storage volume at the end of the pe-  
400 riod. Similar to the forecast for 2017, the 2018 forecast fails to identify when the RSV  
401 will first increase, along with when the storage volume will peak. On the other hand, the  
402 forecast maintains an accurate prediction for the final expected volume in the reservoir.  
403 The observed runoff volumes for all other test years were in excess of 28,000 ac-ft, whereas  
404 2018 demonstrated a runoff volume on the low end of 18,000 ac-ft.

405       Results from the other statistical models are shown in Figure 7. The graphs rep-  
406 resent the monthly storage volume forecasts for all three models across the five consec-  
407 utive test years. SARIMA and TBATS exhibited higher accuracy in capturing seasonal  
408 fluctuations than VAR, as shown respectively in 2015, 2017, and 2019; and 2016, 2018,  
409 and 2019. VAR generally demonstrated the lowest accuracy but occasionally excelled in  
410 forecasting extreme inflow events (e.g. in 2017 and 2019). Overall, ResCNN-LSTM out-  
411 performs the statistical methods except for the 2018 test year (Figures 6 and 7, and Ta-  
412 ble 2). The ResCNN-LSTM produces accurate PSV forecasts for each test year, which  
413 lie within the distribution bounds of the batch forecasts (Figure 8, comparing the PSV  
414 for each test year in comparison to the observed amount). The distributions for the PSV  
415 are skewed upwards towards the reservoir spill limit during each test year. Table 3 sum-  
416 marizes the absolute percent errors at the edges and center of the confidence interval for  
417 each test year, along with forecasts from the statistically-based models. With the excep-

418 tion of SARIMA in 2016 and VAR in 2019, the ResCNN-LSTM consistently outperforms  
419 all three statistical models when using the upper bound of the confidence interval.

## 420 **4 Discussion**

421 While the ResCNN-LSTM approach exhibited higher accuracy overall, there was  
422 significant inter-annual variability, which appears to relate to SWE characteristics. A  
423 relatively low snowpack was observed during the 2014-2015 winter and spring seasons  
424 (Figure 3). The buildup of SWE begins early in October and peaks in the middle of March.  
425 The improved forecast for 2016 may be attributed to a more typical winter experienced  
426 from the SWE monitoring sites. The forecast most similar to 2016 is 2019 in which the  
427 times series for SWE both steadily climb through the winter into the spring season. The  
428 buildup of SWE for 2017 is different from other years: it appears to peak early in March  
429 but then continues to linger through April until finally dropping off at the start of May.  
430 In comparison to the winters experienced in 2016 and 2019, the buildup of SWE peaks  
431 at the start and middle of April, respectively. The most accurate test periods (2015, 2016,  
432 2019) are likely a result of the accumulated SWE tapering off between March and April.  
433 These periods occur at either the tail-end or just outside of the input window, making  
434 them close enough for the model to effectively map the resulting runoff period assum-  
435 ing that no further SWE accumulation takes place. Thus, the timing of snow ablation  
436 dictates model performance as it marks the transition period from SWE accumulation  
437 to runoff fed by snow melt. The 2017 test period is a testament to this theory as the ac-  
438 cumulation of SWE lingers well past April resulting in very poor regression metrics.

439 The proposed model's forecasts are also speculated to be influenced by non-linear  
440 relationships between RSV (Figure 2) and SWE (Figure 3). Similar to 2015, the 2018  
441 test year experienced a winter with below average SWE accumulation. Likewise, both  
442 2015 and 2018 test years have peak SWE at the middle and end of March, respectively.  
443 However, the total observed reservoir inflows are significantly different with 29,287 ac-  
444 ft in 2015 and 18,590 ac-ft in 2018. This indicates the model's learned ability to map non-  
445 linear relationships among the multivariate dataset which statistical methods fail to cap-  
446 ture (e.g. the 2018 VAR forecast, which exhibits low accuracy). Moreover, the model's  
447 best long-term forecasts (Figure 6) were during years of large runoff, making it most ef-  
448 fective during years when accurate prediction is most important. The model's high pre-  
449 dictive performance is likely associated with these events because they historically oc-

450 cur most frequently. Across the entire 30-year dataset for RSV, the most frequently ob-  
451 served reservoir inflow volumes (taken between first of April and beginning of July) were  
452 in excess of 28,000 ac-ft. The ResCNN-LSTM is a data-driven model designed to remem-  
453 ber essential features over time; therefore, this behavior is entirely expected.

454 In addition to inter-annual dependencies, forecasts also depend on reservoir oper-  
455 ations, hydrological features, and other structural characteristics of the water system (Anghileri  
456 et al., 2016). All of these components factor into long-term decision making by water  
457 managers who depend on reliable estimates for peak inflows. The proposed ResCNN-  
458 LSTM proves to effectively forecast the PSV at Upper Stillwater by (i) taking into ac-  
459 count the spill limit as a physical limiter, and (ii) improving accuracy through statisti-  
460 cal confidence intervals. The ResCNN-LSTM exhibits a higher bias in forecasting the  
461 PSV by consistently producing a realistic value close to but never exceeding what the  
462 reservoir can physically store (Figure 6). In contrast to the statistical models (Figure  
463 7), the forecasted PSV is widely over-predicted during some years (e.g. SARIMA: 2015  
464 and 2018; VAR: 2015) and under-predicted during others (e.g. TBATS: 2015 - 2017 and  
465 2019; VAR: 2016 and 2019). This behavior is most likely attributed to the fact that Up-  
466 per Stillwater has a spill limit to account for storage capacity. Thus, the spill limit acts  
467 as an asymptote for the model when producing each forecast, whereas other statistical  
468 models fail to recognize it. Likewise, this physical feature would account for the upwards  
469 skew in distribution for each test year’s PSV forecast (Figure 8). The confidence inter-  
470 val also provides valuable insight for future decision making. The most challenging year  
471 for the model was 2018 due to an abnormally low snowpack throughout the winter. With  
472 this taken into consideration, the model’s forecast for PSV in 2018 still lies within a phys-  
473 ically realistic value and statistically confident interval. With an average absolute per-  
474 cent error of 2.66% in the center and 1.82% in the upper bound (Table 3), the confidence  
475 interval’s upper bound consecutively proves to be the most accurate in predicting the  
476 PSV for all five test years. Therefore, water managers would be best served by using the  
477 forecasted upper bound for PSV in their forecasts.

478 Geared towards strengthening water manager abilities to manage and conserve RSV,  
479 this research improves on prior multi-step reservoir forecasting efforts by effectively in-  
480 creasing the context size to capture temporal dependencies. Previous studies including  
481 direct-step (Coulibaly et al., 2005; Sattari et al., 2012; Bai et al., 2016) and multi-step  
482 (Coulibaly et al., 2000; Muluye & Coulibaly, 2007; Kao et al., 2020) deep-learning al-

483 algorithms have improved forecasting accuracy, but have yet to accurately forecast peak  
484 inflows at extended long-term horizons. For a lead time of one week, Coulibaly et al. (2000)  
485 forecasted peak flows ranging from an underprediction of 4.3% to an overprediction of  
486 3.8% on average. Similarly, extended four month forecasts by Muluye and Coulibaly (2007)  
487 demonstrated reasonable predictions of low and medium reservoir inflows, but then ei-  
488 ther under or over-predicted the peaks. By contrast, our proposed model consistently  
489 manages to confidently predict realistic PSV values at a three-month lead time with an  
490 average absolute percent error of 2.66%. Similar to Coulibaly et al. (2000), the long-term  
491 forecasts at the end of the forecasted runoff period were more accurate than the initial  
492 short-term values. Periods where the proposed model struggled the most are irregular  
493 runoff seasons with abnormally dry hydrologic conditions and late-season SWE accumu-  
494 lation. An extended horizon of three months leaves wide potential for changes in hydro-  
495 logic conditions and can amount to multiple different scenarios for runoff. A topic for  
496 future research may involve the inclusion of scenario-based model runs incorporating me-  
497 teorological predictions from outside entities.

## 498 **5 Conclusion**

499 Given the comparatively high performance of the proposed algorithm in the study  
500 region, the ResCNN-LSTM architecture warrants further study for multi-step RSV fore-  
501 casting. Considerations for future research include (i) further experiments with model  
502 architecture, (ii) investigating additional independent variables, and (iii) modeling ad-  
503 ditional reservoirs influenced by snowmelt runoff. Potential experiments with model ar-  
504 chitecture may include implementing batch normalization (Ioffe & Szegedy, 2015) be-  
505 tween layers to reduce training time and increase predictive accuracy, and utilizing an  
506 attention mechanism (Bahdanau et al., 2014) to observe the intermediate states of the  
507 encoder, rather than only the final states. The inclusion of other independent variables,  
508 such as atmospheric temperature and solar radiation, may further improve accuracy. Fi-  
509 nally, modeling additional reservoirs will provide valuable insight into model transfer-  
510 ability.

## 511 **Acknowledgments**

512 This research was supported by the Central Utah Water Conservancy District and Uni-  
513 versity of Utah. A special thanks to Blake Buehler, Heath Clark, Jared Hansen, and Zac

514 Brekenridge for their useful comments and discussion. Geographic Information System  
515 (GIS) mapping was provided by Lindsay Bentley. All reservoir storage volume (RSV) and  
516 snow water equivalent (SWE) data is made publicly available by the National Resource  
517 Conservation Service online repository at <https://www.nrcs.usda.gov/>. The Python code  
518 and training data, including the scripts used to compile them, are available through figshare  
519 (10.6084/m9.figshare.12756014).

## 520 References

- 521 Agarap, A. F. (2018). Deep learning using rectified linear units (relu). *arXiv*  
522 *preprint arXiv:1803.08375*.
- 523 Ahmed, S., Coulibaly, P., & Tsanis, I. (2015). Improved spring peak-flow forecasting  
524 using ensemble meteorological predictions. *Journal of Hydrologic Engineering*,  
525 *20*(2), 04014044.
- 526 Anghileri, D., Voisin, N., Castelletti, A., Pianosi, F., Nijssen, B., & Lettenmaier,  
527 D. P. (2016). Value of long-term streamflow forecasts to reservoir opera-  
528 tions for water supply in snow-dominated river catchments. *Water Resources*  
529 *Research*, *52*(6), 4209–4225.
- 530 Bahdanau, D., Cho, K., & Bengio, Y. (2014). *Neural machine translation by jointly*  
531 *learning to align and translate*.
- 532 Bai, Y., Chen, Z., Xie, J., & Li, C. (2016). Daily reservoir inflow forecasting using  
533 multiscale deep feature learning with hybrid models. *Journal of Hydrology*,  
534 *532*, 193 - 206.
- 535 Bai, Y., Wang, P., Xie, J., Li, J., & Li, C. (2015). Additive model for monthly reser-  
536 voir inflow forecast. *Journal of Hydrologic Engineering*, *20*(7), 04014079.
- 537 Bennett, J. C., Wang, Q. J., Li, M., Robertson, D. E., & Schepen, A. (2016). Reli-  
538 able long-range ensemble streamflow forecasts: Combining calibrated climate  
539 forecasts with a conceptual runoff model and a staged error model. *Water*  
540 *Resources Research*, *52*(10), 8238-8259.
- 541 Bourdin, D. R., Nipen, T. N., & Stull, R. B. (2014). Reliable probabilistic fore-  
542 casts from an ensemble reservoir inflow forecasting system. *Water Resources*  
543 *Research*, *50*(4), 3108-3130.
- 544 Box, G. E., Jenkins, G. M., Reinsel, G. C., & Ljung, G. M. (2015). *Time series*  
545 *analysis: forecasting and control*. John Wiley & Sons.

- 546 Budu, K. (2014). Comparison of wavelet-based ann and regression models for reser-  
547 voir inflow forecasting. *Journal of Hydrologic Engineering*, 19(7), 1385-1400.
- 548 Cengil, E., & Cinar, A. (2018). A deep learning based approach to lung cancer iden-  
549 tification. , 1-5.
- 550 Chang, F.-J., & Tsai, M.-J. (2016). A nonlinear spatio-temporal lumping of radar  
551 rainfall for modeling multi-step-ahead inflow forecasts by data-driven tech-  
552 niques. *Journal of Hydrology*, 535, 256 - 269.
- 553 Chiamsathit, C., Adeloye, A. J., & Bankaru-Swamy, S. (2016). Inflow forecasting us-  
554 ing artificial neural networks for reservoir operation. *Proceedings of the Inter-  
555 national Association of Hydrological Sciences*, 373, 209-214.
- 556 Coulibaly, P., Anctil, F., & Bobée, B. (2000). Daily reservoir inflow forecasting us-  
557 ing artificial neural networks with stopped training approach. *Journal of Hy-  
558 drology*, 230(3), 244 - 257.
- 559 Coulibaly, P., Haché, M., Fortin, V., & Bobée, B. (2005). Improving daily reser-  
560 voir inflow forecasts with model combination. *Journal of Hydrologic Engineer-  
561 ing*, 10(2), 91-99.
- 562 De Livera, A. M., Hyndman, R. J., & Snyder, R. D. (2011). Forecasting time se-  
563 ries with complex seasonal patterns using exponential smoothing. *Journal of  
564 the American statistical association*, 106(496), 1513-1527.
- 565 Elizaga, N. B., Maravillas, E. A., & Gerardo, B. D. (2014). Regression-based inflow  
566 forecasting model using exponential smoothing time series and backpropaga-  
567 tion methods for angat dam. In *2014 international conference on humanoid,  
568 nanotechnology, information technology, communication and control, environ-  
569 ment and management (hnicem)* (pp. 1-6).
- 570 Gehring, J., Auli, M., Grangier, D., Yarats, D., & Dauphin, Y. N. (2017). Convo-  
571 lutional sequence to sequence learning. In *Proceedings of the 34th international  
572 conference on machine learning-volume 70* (pp. 1243-1252).
- 573 Gulli, A., & Pal, S. (2017). *Deep learning with keras*. Packt Publishing Ltd.
- 574 He, K., Zhang, X., Ren, S., & Sun, J. (2015). *Deep residual learning for image recog-  
575 nition*.
- 576 Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural compu-  
577 tation*, 9(8), 1735-1780.
- 578 Hyndman, R. J. (2008). *Forecasting with exponential smoothing : the state space ap-*

- 579            *proach*.
- 580    Iddrisu, W., Nokoe, K. S., & Osei-Akoto, I. (2016, 09). Modelling the trend of flows  
581            with respect to rainfall variability using vector autoregression. *International*  
582            *Journal of Advanced Research*, 4, 125-140.
- 583    Ioffe, S., & Szegedy, C. (2015). *Batch normalization: Accelerating deep network*  
584            *training by reducing internal covariate shift*.
- 585    Ivanov, V., & Kilian, L. (2005). A practitioner's guide to lag order selection for  
586            var impulse response analysis. *Studies in Nonlinear Dynamics Econometrics*,  
587            9(1).
- 588    Kao, I.-F., Zhou, Y., Chang, L.-C., & Chang, F.-J. (2020). Exploring a long short-  
589            term memory based encoder-decoder framework for multi-step-ahead flood  
590            forecasting. *Journal of Hydrology*, 124631.
- 591    Kingma, D. P., & Ba, J. (2014). Adam: A method for stochastic optimization. *arXiv*  
592            *preprint arXiv:1412.6980*.
- 593    Konapala, G., & Mishra, A. (2020). Quantifying climate and catchment control  
594            on hydrological drought in the continental united states. *Water Resources Re-*  
595            *search*, 56(1), e2018WR024620.
- 596    Kote, A. S., & Jothiprakash, V. (2009). Monthly reservoir inflow modeling using  
597            time lagged recurrent networks. *International Journal of Tomography and*  
598            *Statistics*, 12(F09), 64–84.
- 599    Kratzert, F., Klotz, D., Herrnegger, M., Sampson, A. K., Hochreiter, S., & Nearing,  
600            G. S. (2019). Toward improved predictions in ungauged basins: Exploiting the  
601            power of machine learning. *Water Resources Research*.
- 602    LeCun, Y. e. a. (1990). Handwritten digit recognition with a back-propagation net-  
603            work. *Advances in Neural Information Processing Systems*, 396–404.
- 604    Liu, H., & Song, B. (2018). Stock price trend prediction model based on deep resid-  
605            ual network and stock price graph. , 02, 328-331.
- 606    Liu, Y., Sang, Y.-F., Li, X., Hu, J., & Liang, K. (2017). Long-term streamflow fore-  
607            casting based on relevance vector machine model. *Water*, 9(1).
- 608    Muluye, G. Y., & Coulibaly, P. (2007). Seasonal reservoir inflow forecasting with  
609            low-frequency climatic indices: a comparison of data-driven methods. *Hydro-*  
610            *logical sciences journal*, 52(3), 508–522.
- 611    Ning, D., Liu, G., Jiang, R., & Wang, C. (2019). Attention-based multi-scale trans-

- 612           fer resnet for skull fracture image classification. , 11198, 111980D–111980D-5.
- 613       NRCS. (2020 (accessed May 3, 2020)). Nrcs online api [Computer software manual].
- 614           Retrieved from <https://www.nrcs.usda.gov/>"
- 615       Papamichail, D. M., & Georgiou, P. E. (2001). Seasonal arima inflow models for
- 616           reservoir sizing 1. *JAWRA Journal of the American Water Resources Associa-*
- 617           *tion*, 37(4), 877–885.
- 618       Robbins, H., & Monro, S. (1951). A stochastic approximation method. *The annals*
- 619           *of mathematical statistics*, 400–407.
- 620       Sattari, M. T., Yurekli, K., & Pal, M. (2012). Performance evaluation of artificial
- 621           neural network approaches in forecasting reservoir inflow. *Applied Mathemati-*
- 622           *cal Modelling*, 36(6), 2649 - 2657.
- 623       Shumway, R. H., & Stoffer, D. S. (2000). Time series analysis and its applications.
- 624           *Studies In Informatics And Control*, 9(4), 375–376.
- 625       Sims, C. A. (1980). Macroeconomics and reality. *Econometrica: journal of the*
- 626           *Econometric Society*, 1–48.
- 627       Sin, C.-Y., & White, H. (1996). Information criteria for selecting possibly misspeci-
- 628           fied parametric models. *Journal of Econometrics*, 71(1), 207–225.
- 629       Stokelj, T., Paravan, D., & Golob, R. (2002). Enhanced artificial neural network in-
- 630           flow forecasting algorithm for run-of-river hydropower plants. *Journal of Water*
- 631           *Resources Planning and Management*, 128(6), 415–423.
- 632       Sutskever, I., Martens, J., Dahl, G., & Hinton, G. (2013). On the importance of
- 633           initialization and momentum in deep learning. In *International conference on*
- 634           *machine learning* (pp. 1139–1147).
- 635       Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with
- 636           neural networks. In *Proceedings of the 27th international conference on neural*
- 637           *information processing systems - volume 2* (pp. 3104–3112). Cambridge, MA,
- 638           USA: MIT Press.
- 639       Tukey, J. W. (1970). *Exploratory data analysis*. Addison-Wesley.
- 640       van den Oord, A., Dieleman, S., Zen, H., Simonyan, K., Vinyals, O., Graves, A., ...
- 641           Kavukcuoglu, K. (2016). *Wavenet: A generative model for raw audio*.
- 642       Wang, F., Wang, L., Zhou, H., Saavedra Valeriano, O. C., Koike, T., & Li, W.
- 643           (2012). Ensemble hydrological prediction-based real-time optimization of a
- 644           multiobjective reservoir during flood season in a semiarid basin with global

- 645 numerical weather predictions. *Water Resources Research*, 48(7).
- 646 Yen, C.-F., Hsieh, H.-Y., Su, K.-W., & Leu, J.-S. (2019). Predicting solar perfor-  
 647 mance ratio based on encoder-decoder neural network model. In *2019 11th*  
 648 *international congress on ultra modern telecommunications and control systems*  
 649 *and workshops (icumt)* (pp. 1–4).
- 650 Yuan, M., Ji, X., Lu, T., Chen, P., & Zhang, H. (2019). A novel two-factor attention  
 651 encoder-decoder network through combining temporal and prior knowledge for  
 652 weather forecasting. In *2019 international joint conference on neural networks*  
 653 *(ijcnn)* (pp. 1–8).
- 654 Zhang, X., & You, J. (2020). A gated dilated causal convolution based encoder-  
 655 decoder for network traffic forecasting. *IEEE Access*, 8, 6087–6097.
- 656 Zheng, Z., Molotch, N. P., Oroza, C. A., Conklin, M. H., & Bales, R. C. (2018).  
 657 Spatial snow water equivalent estimation for mountainous areas using wireless-  
 658 sensor networks and remote-sensing products. *Remote sensing of environment*,  
 659 215, 44–56.
- 660 Zhou, Y., Guo, S., & Chang, F.-J. (2019). Explore an evolutionary recurrent anfis  
 661 for modelling multi-step-ahead flood forecasts. *Journal of Hydrology*, 570, 343  
 662 - 355.

## 663 6 Tables

**Table 1.** Optimal architecture and configuration for the model

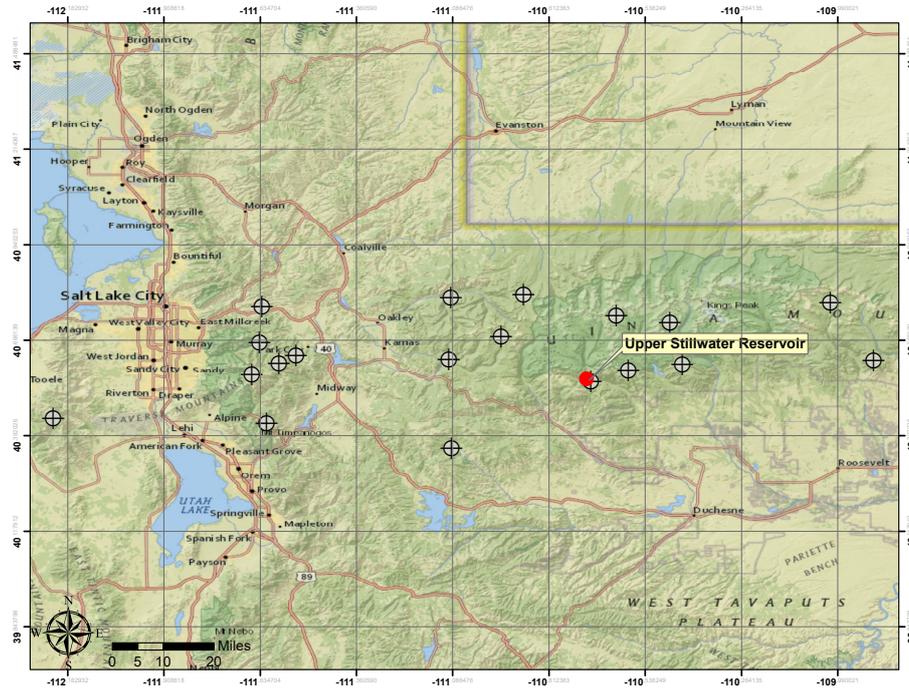
Parameter	Selected Value	Tested Values
Input Window	20	15,20,25
Epochs	100	50,75,100,150
Batch Size	32	16,32,64
Filter1 (CNN)	16	16,32,64
Filter2 (CNN)	32	16,32,64
Filter3 (CNN)	64	16,32,64
Kernel (CNN)	6	4,5,6,7,8
Cells (LSTM)	64	16,32,64,128
LSTM Layers	4	1,2,3,4,5
CNN Layers	12	3,6,12,18

**Table 2.** Regression metrics for ResCNN-LSTM hold-out sets

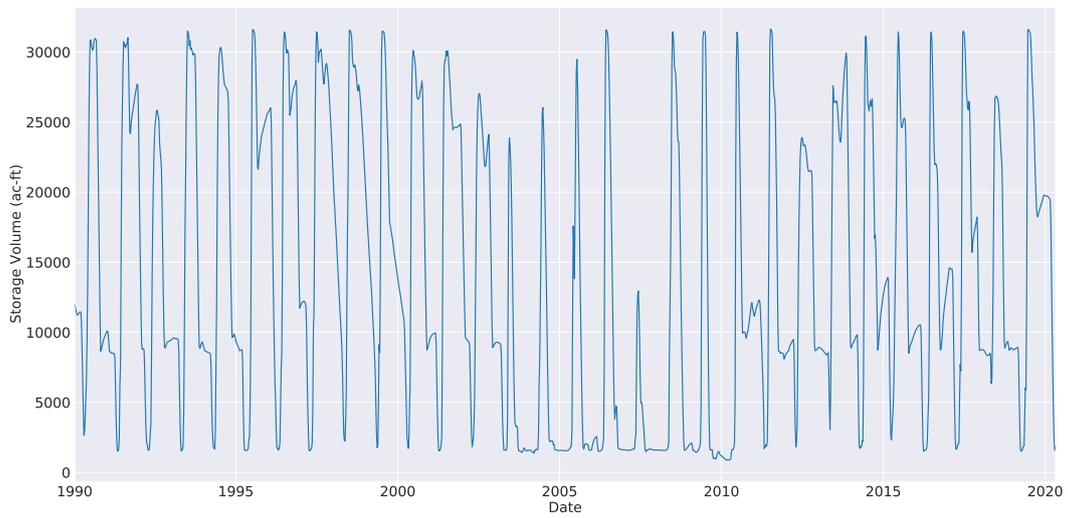
	MAE (%)	RMSE (%)	MedAE (%)	NSE	ExpVar
2015	22.949	28.090	19.161	0.889	0.932
2016	12.295	16.695	8.132	0.968	0.972
2017	32.863	43.216	29.509	0.792	0.836
2018	60.759	70.335	48.331	0.412	0.851
2019	21.515	35.361	7.208	0.812	0.862

**Table 3.** Forecasted PSV absolute percent error

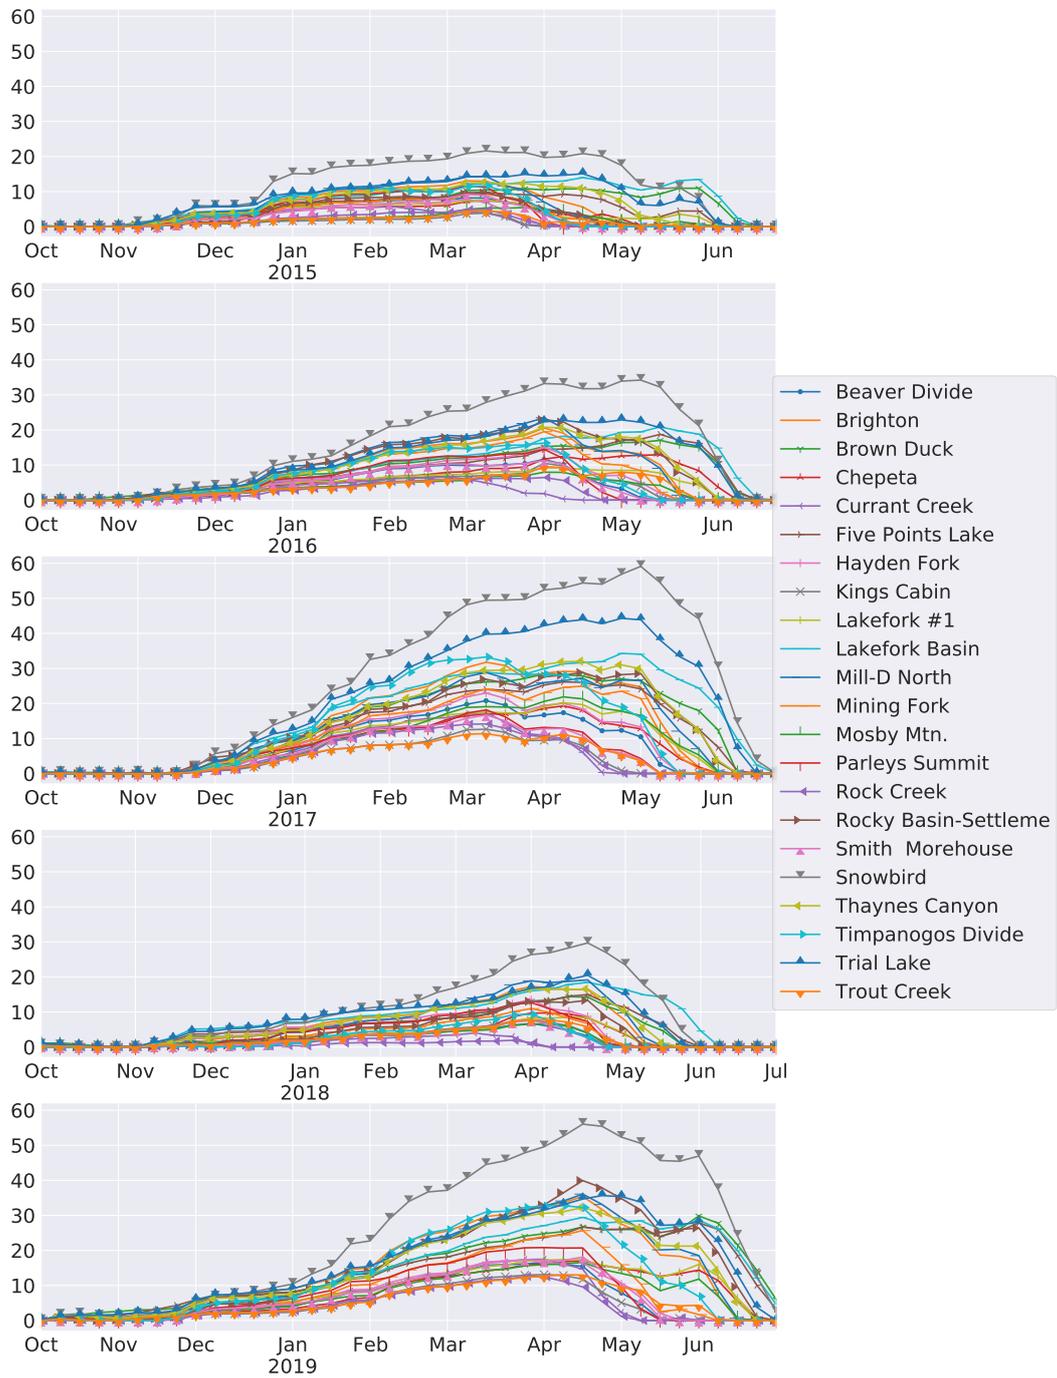
	Lower Limit	Center	Upper Limit	SARIMA	VAR	TBATS
2015	2.979	2.136	1.294	11.874	14.343	14.144
2016	3.529	2.703	1.878	1.526	28.256	16.396
2017	6.029	4.195	2.362	4.792	3.452	17.607
2018	7.020	2.301	2.419	47.989	43.745	12.602
2019	2.858	1.995	1.133	3.830	0.924	8.356



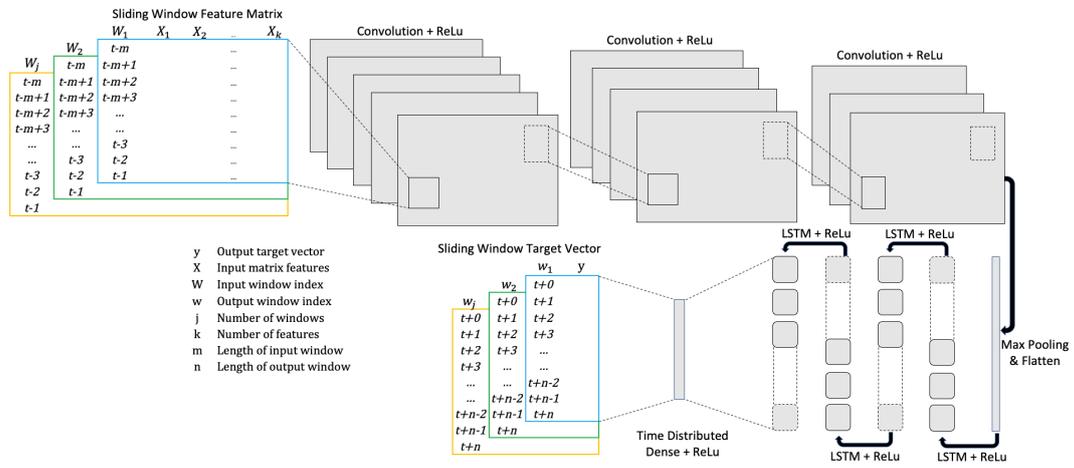
**Figure 1.** Study site location: Upper Stillwater reservoir, Utah



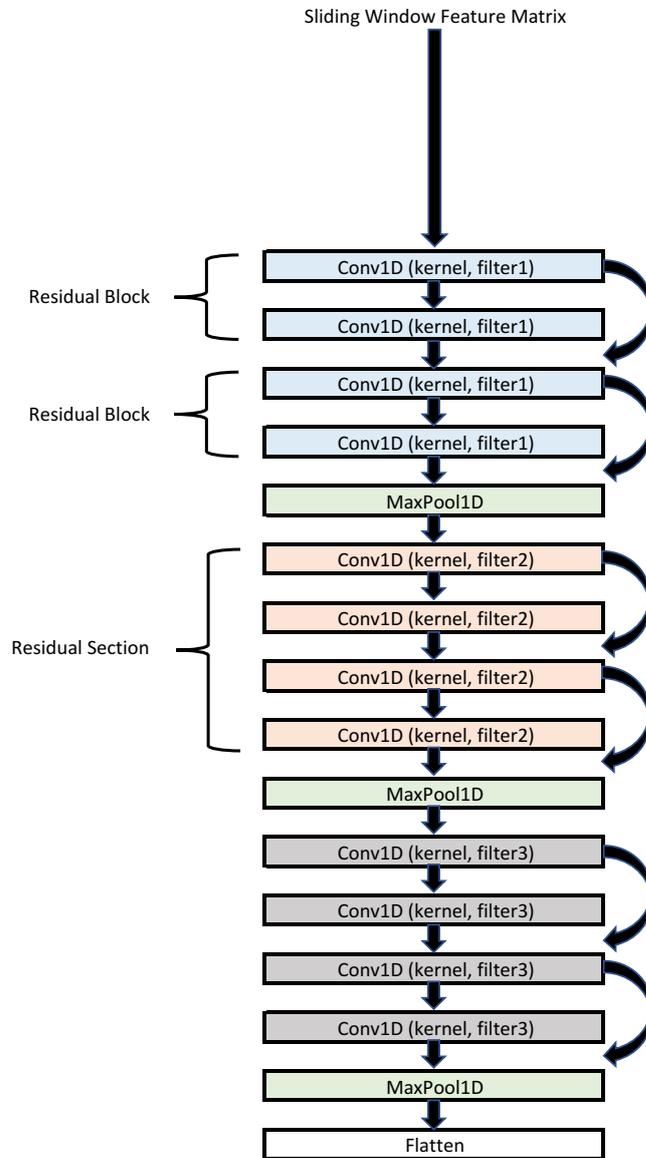
**Figure 2.** Upper Stillwater historical RSV (ac-ft)



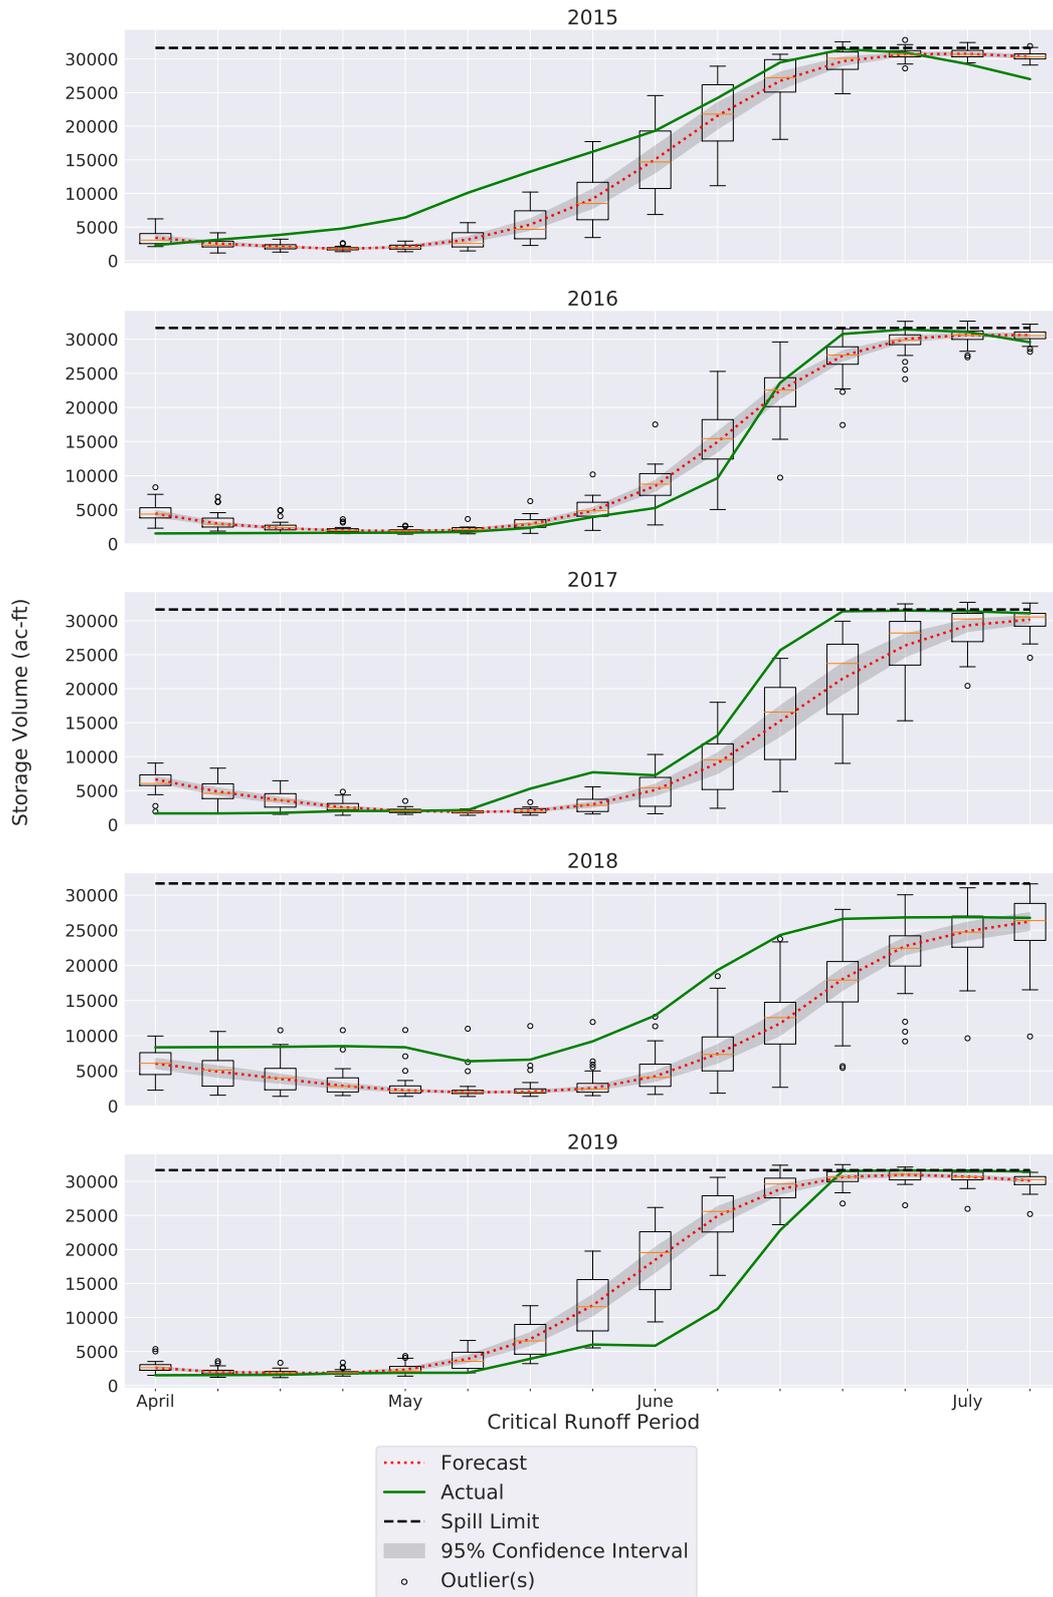
**Figure 3.** Recorded SWE (in) model inputs for five consecutive test years



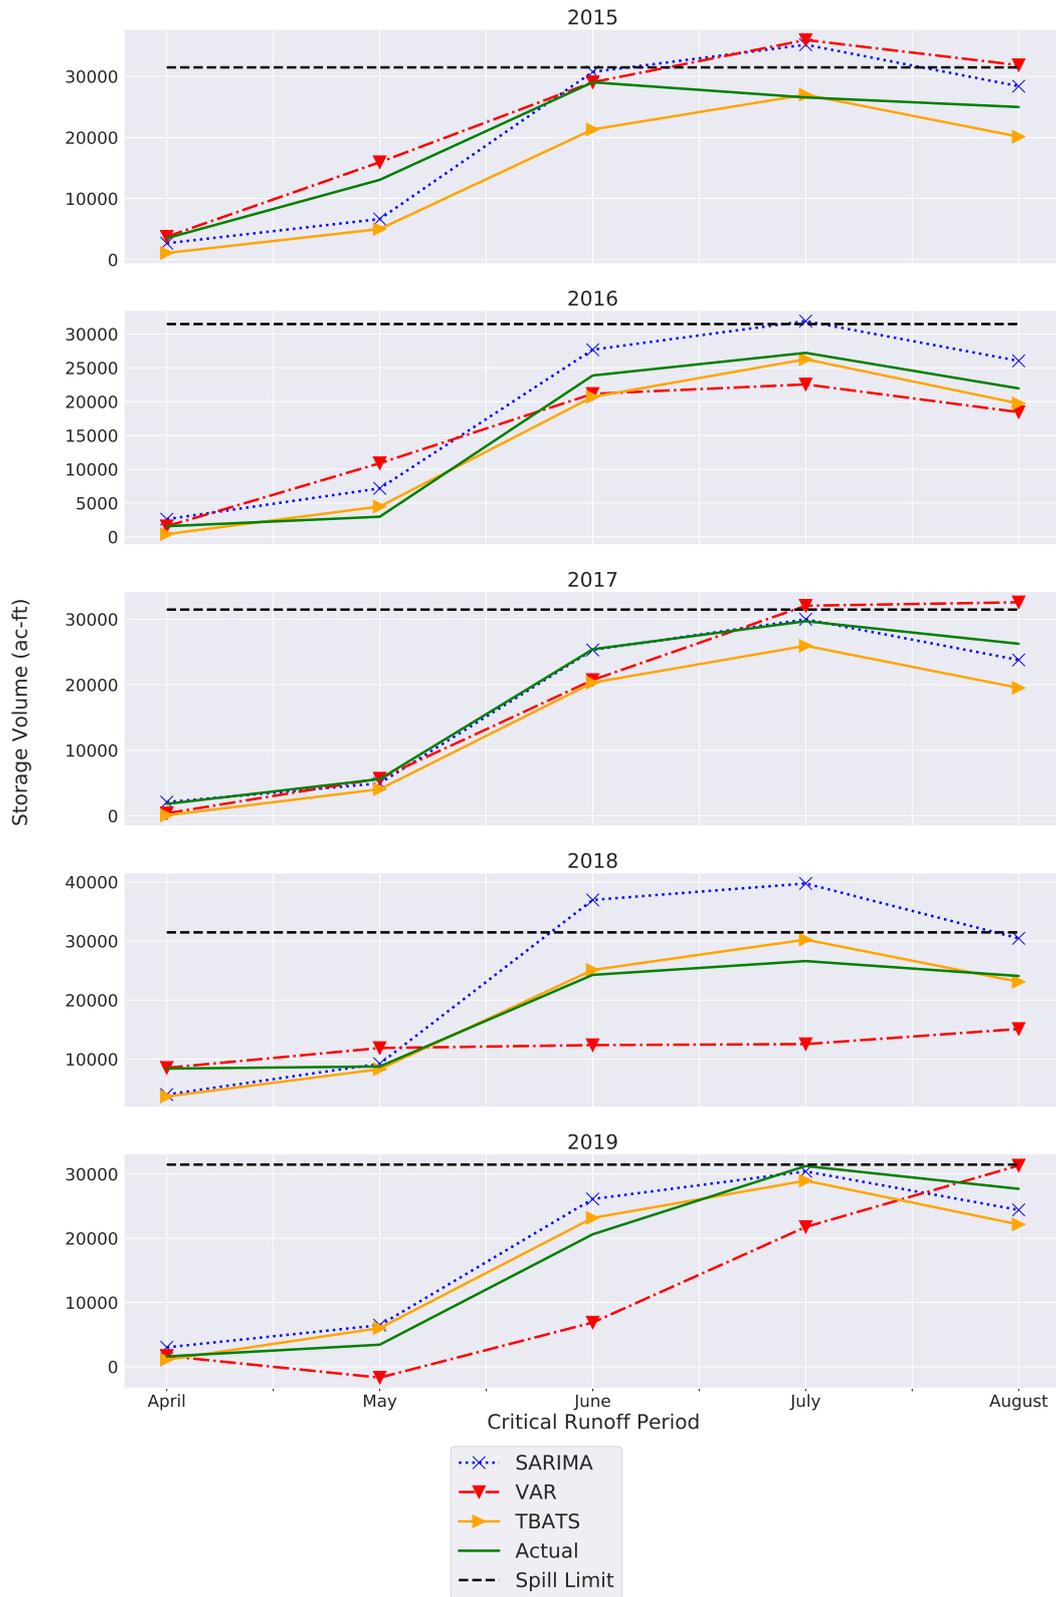
**Figure 4.** Deep learning model architecture



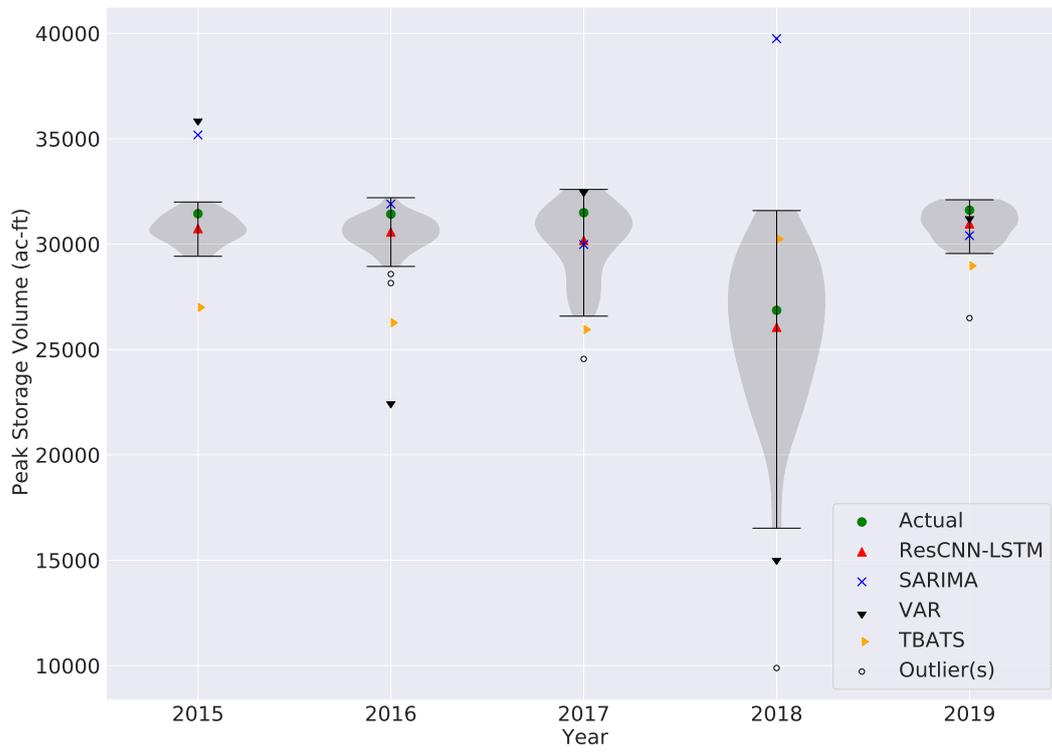
**Figure 5.** Deep residual framework for CNN encoder



**Figure 6.** Multi-step storage volume forecasts from 2015 through 2019 using ResCNN-LSTM following a weekly timestep frequency



**Figure 7.** Multi-step storage volume forecasts from 2015 through 2019 using classic statistical methods following a monthly timestep frequency



**Figure 8.** Comparison of various model's forecasted PSV