PROCESS-BASED ASSIMILATION OF REMOTELY SENSED SNOW DATA INTO A DISTRIBUTED PHYSICAL HYDROLOGICAL MODEL

Elijah N. Boardman[1]

ABSTRACT

Process-based assimilation of remotely sensed snow data into a distributed physical hydrological model can reduce error in seasonal water supply forecasts. Snow data assimilation corrects errors in the modeled volume and spatial distribution of snow water equivalent (SWE), which arise from precipitation data errors, rain / snow partitioning errors, and ablation dynamics errors. However, data assimilation in physical models also mandates comprehensive reanalysis of the model's history and current state to maintain internal consistency and avoid violations of the water mass balance. The method introduced here uses a spatially explicit approximation of the pre-assimilation model dynamics to construct a time-reversible surrogate model, from which errors in precipitation, streamflow, and storage states can all be corrected based on a single SWE map. An application of process-based assimilation is demonstrated with the Distributed Hydrology Soil Vegetation Model (DHSVM) using SWE maps from the Airborne Snow Observatories (ASO). In an 11-year backcast simulation, DHSVM with process-based SWE assimilation achieves a mean absolute percent error of 7% for April-July inflow to Hetch Hetchy Reservoir, compared to 27% error without assimilation and 19% error with direct insertion. The union of model and data through process-based assimilation is critical to realizing the value of SWE maps.
(KEYWORDS: water supply forecast, physical modeling, remote sensing, data assimilation)

INTRODUCTION

Seasonal water supply forecasts are a core component of water management planning in the western U.S., where cold-season snowpack measurements can help predict warm-season runoff (Church 1933, Pagano et al. 2004). Forecasting is inherently an inexact science, and reducing uncertainty remains a priority (Schaake and Peck 1985, Stillinger et al. 2021). Considerable effort has been invested in improved quantification of the snowpack through remote sensing (Nolin 2010). The ASO program is one of the fastest growing approaches to snow remote sensing, using airborne lidar surveys of snow depth with modeled and measurement-constrained density fields to produce full-basin SWE maps (Painter et al. 2016). While accurate snowpack quantification is a crucial first step to improve water supply forecasts, snow is only one component of the mountain water balance. Water supply forecasts must also account for fluxes such as rain, interception, sublimation, evapotranspiration, and changes in groundwater or soil moisture, in addition to controls on runoff timing such as snowmelt and flow routing. In order to build more resilient forecast methods that can benefit from spatial snow data and account for climate nonstationarity and disturbance, recent efforts have operationalized distributed physical hydrological models for seasonal water supply forecasting (e.g., WRF-Hydro, Gochis et al. 2020, and DHSVM-WSF, Boardman 2023). The best way to combine advances in snow remote sensing with advances in process-based hydrological modeling is an open question.

Assimilation by direct insertion works for snow models but is invalid for physical hydrological models. Direct insertion refers to running a model up to a snow measurement date (e.g., an ASO flight), replacing modeled snow data with remotely sensed snow data, and resuming the model run. Direct insertion of ASO data into the iSnobal snow model significantly improves the snowpack simulation (Hedrick et al. 2018), but models that only simulate snowpack dynamics are fundamentally different from models that simulate the entire mountain water cycle. Snow models can adjust to direct insertion because they only need to account for processes affecting the snowpack. Hydrological models like DHSVM (Wigmosta et al. 1994) are unsuitable for direct insertion because the updated snow state would be inconsistent with non-snow components of the model. For example, if the modeled snowpack is too small because the model incorrectly partitioned prior precipitation as rain instead of snow or melted the snowpack too quickly, direct insertion of a deeper snowpack will cause a mass balance issue wherein the same precipitation is double counted in the model. From a mass balance perspective, direct insertion implies that all snowpack errors are due to precipitation errors, but this is a false premise because errors in the modeled snowpack evolution between ASO flights are observed even when no precipitation occurred. This paper poses the question: what steps are necessary to more robustly assimilate SWE maps into distributed physical hydrological models?

Paper presented Western Snow Conference 2024
[1]Mountain Hydrology LLC, Reno, NV eli.boardman@mountainhydrology.com

ASSIMILATION METHOD

This paper presents a process-based snow data assimilation method that preserves the benefits of direct insertion (namely, exactly matching the measured SWE map) while also preserving the water mass balance and maintaining internal consistency in the model's non-snow process representations. At its most basic, this is accomplished by changing non-snow water balance components so they maintain consistency with the updated snow state. In practice, this is rather tricky, because distributed process-based hydrological models are akin to "Rube Goldberg machines" wherein everything is connected to everything else, so localized changes can have complex and nonlinear downstream effects.

Causes of Snowpack Error

Errors in snowpack simulations can arise from numerous sources, not just precipitation data errors. If all snowpack errors in physical hydrological models were a result of precipitation data errors, it would be valid to implement direct insertion of measurement-based SWE data because any increase or decrease in the total water mass balance would offset the inferred precipitation error. However, errors in the modeled snowpack on a particular date can also result from imperfect simulation of snowpack processes like sublimation and melt, or from rain / snow partitioning errors attributable to meteorological forcing data (i.e., temperature) or model assumptions. Experience suggests that in relatively well-instrumented mountain ranges like the Sierra, most discrepancies between modeled and measured SWE maps can be explained by errors in model dynamics without having to invoke a precipitation data error. Since precipitation forcing data are ostensibly based on at least some measurements, whereas snowpack dynamics are purely simulated, it is preferable to preserve whatever limited information might exist in meteorological forcing data when possible. In other words, blaming "data error" should be the last resort of the hydrological modeler; in most cases, it is the model that is wrong, not the data.

The Danger of Mass Balance Violations

Failure to account for non-snow errors in physical hydrological models makes direct insertion of snow data a dangerous proposition. Suppose a modeled SWE map shows 1.35 km^3 (1,100 TAF) of SWE in a particular watershed, but a recent measurement-based SWE map only shows 1.23 km^3 (1,000 TAF) of SWE on the same date. This magnitude of error (~10%) is not uncommon in snowpack models. Suppose also that the modeled SWE volume error is a result of too little snowmelt in the model, perhaps due to dust deposition causing lower-than-normal snow albedo and increased radiative forcing that is imperfectly represented in the model. The extra SWE present in the model should have been melted earlier, and the resultant water might still be partially stored in the subsurface in addition to contributing to antecedent streamflow. Updating only the model's snowpack state with direct insertion violates the water mass balance, effectively causing 0.12 km^3 (100 TAF) of water to permanently vanish into what can only be interpreted in physical terms as a black hole. As a result, future streamflow simulated by the model will be biased low, because the direct insertion procedure does not account for the contribution of missing or extra SWE to storage and streamflow. Indeed, the author's preliminary tests of SWE assimilation in DHSVM-WSF via direct insertion often increased runoff volume errors, pointing to a clear problem and ultimately leading to the development of the method described hereafter.

Time Irreversibility of Physical Models

Thermodynamics puts constraints on the type of approach that can be used to correct errors resulting from pre-assimilation model dynamics. The hallmark of physically based models is time irreversibility. Any model that simulates natural phenomenon involving energy transfer must be time irreversible due to the Second Law of Thermodynamics. For example, in distributed physical hydrological models, it is impossible to run streamflow generation processes in reverse, because there is an increase in entropy when water mixes from different grid cells. Ideally, a hydrological model could be run in reverse starting from the observed SWE state to determine what forcing conditions and dynamics would have been necessary to produce the correct results, but this is strictly impossible due to time irreversibility. Thus, the goal of process-based assimilation is to construct an *approximate* time-reversible surrogate model, making the minimum assumptions necessary to answer: what would have been different about the pre-assimilation model dynamics if the model had more accurately simulated the snowpack?

Process-Based Approach to Assimilation

Process-based snow data assimilation corrects inferred errors in non-snow states based on pre-assimilation water mass balance fluxes. The model's water mass balance history is used to quantify relationships between precipitation, snow, streamflow, and storage, and these relationships are used to propagate SWE errors into other

mass balance components. Inferred errors in the antecedent streamflow timeseries and current storage state (maps of soil moisture and channel storage) are retroactively corrected before the model run is resumed with the updated SWE state. Figure 1 depicts a cross-section of a distributed physical hydrological model, illustrating the application of process-based snow data assimilation. In the first panel, the model state has substantial groundwater and channel water storage, high antecedent streamflow, and a small snowpack as a result of mixed rain and snowfall. Compared to the measurement-informed SWE map, most grid cells have too little snow, which means that the bulk SWE error is negative, though some cells have too much snow. For every single grid cell, the positive or negative SWE error is propagated downhill through the subsurface, increasing or decreasing the groundwater and soil moisture until reaching a channel. Antecedent streamflow is reduced, and the channel storage is updated for congruency with the most-recent streamflow rate. After iterating through all grid cells, the total mass of water (snow + liquid storage + antecedent streamflow) is unchanged, but this water is distributed differently (in space and in time) so that the updated state and history could plausibly have resulted from accurate simulation of the pre-assimilation snowpack dynamics. In other words, the model is now close to what it *should* have looked like if it had gotten the snow right.

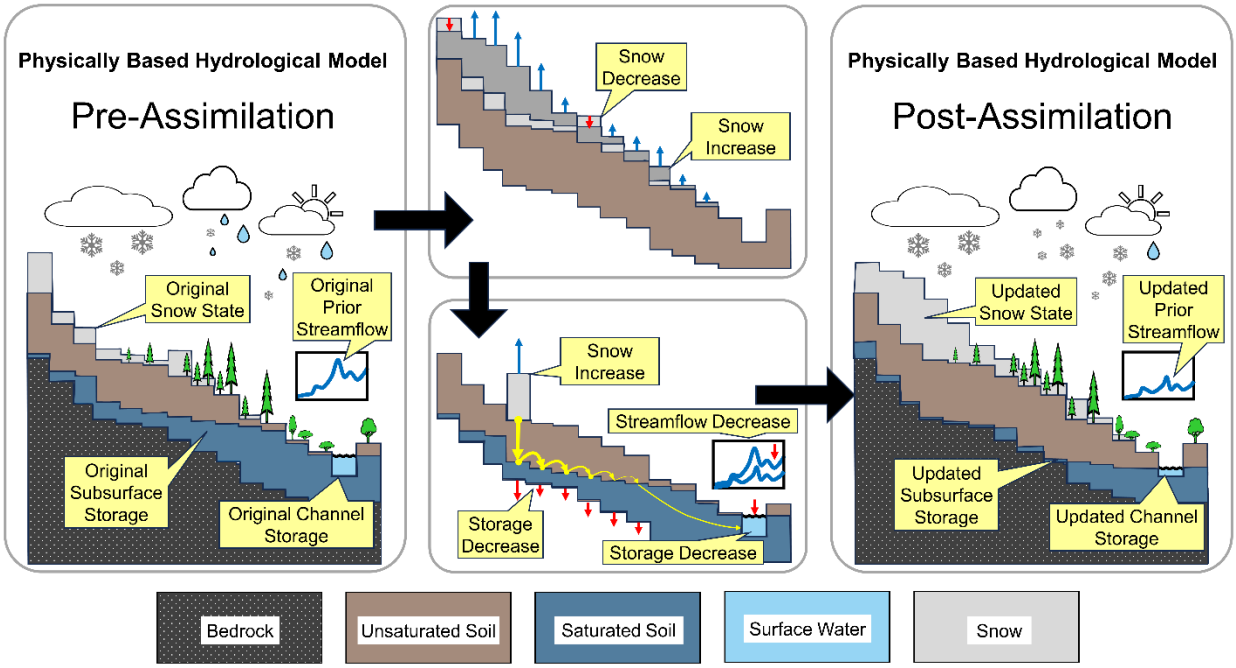

Figure 1. Cross-sectional view of process-based snow data assimilation method applied to a distributed physical hydrological model.

Water Mass Balance Partitioning

During process-based assimilation, the partitioning of water between different fluxes is reanalyzed, but the total water mass is only changed in response to unavoidable precipitation errors. In Figure 2, the total mass balance is represented by initial storage and pre-assimilation precipitation inputs. The model partitions this water between evapotranspiration, streamflow, and solid (snow) or liquid storage. Supposing that the true snowpack is larger than the simulated snowpack, proportional reductions are accorded to antecedent streamflow and storage fluxes (final – initial storage). Sometimes it is necessary to also invoke a precipitation error if the total mass balance cannot be explained by the unaltered meteorological records. The updated model state sums to the original volume (plus or minus any unavoidable precipitation error). In Figure 3, pie charts illustrate the re-partitioning of water mass balance fluxes after snow is observed to be a larger-than-modeled component of the mass balance. Any precipitation error, if unavoidable, increases or decreases the size of each pie chart but does not alter the partitioning. Evapotranspiration is unaltered due to an energy-limited assumption when a snowpack is present. Errors in the partitioning of water between storage and streamflow are assumed to be proportional to the magnitude of each of these fluxes; that is, if 90% of snowmelt went to streamflow on the pre-assimilation period, 90% of the snowpack error will be accounted for by increasing or decreasing streamflow. This assumption of proportionality is a simplification of nonlinear hydrological dynamics, but due to the time irreversibility problem, a better assumption is not forthcoming. After assimilation, the model is internally consistent, physically plausible, and in agreement with measured snow data.

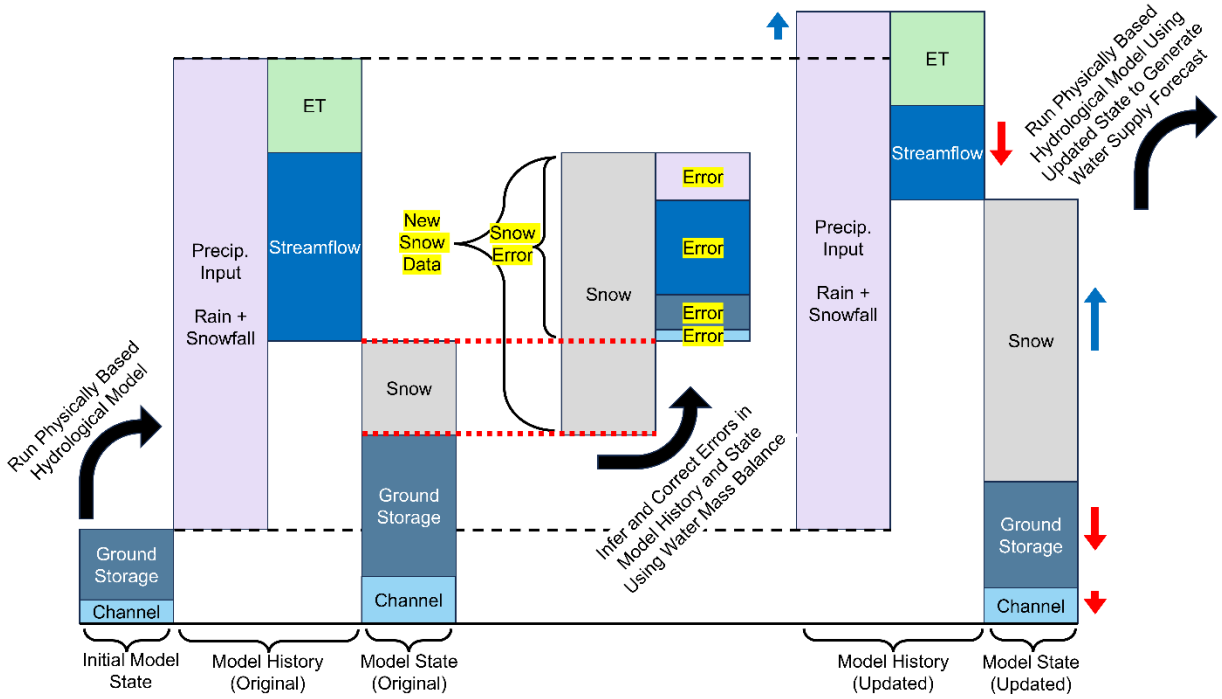

Figure 2. Water mass balance preservation during process-based snow data assimilation.

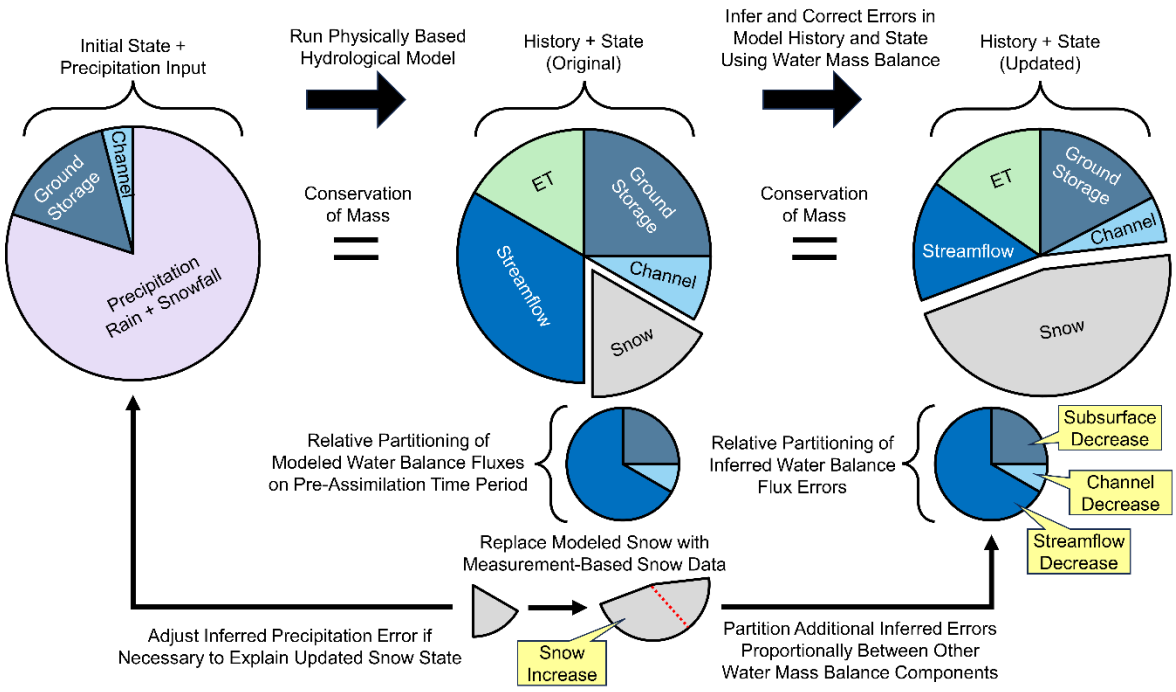

Figure 3. Re-partitioning of water mass balance during process-based snow data assimilation.

Algorithm

The following steps detail the implementation of process-based snow data assimilation for a distributed physical hydrological model. It is assumed that a full-basin measurement-based SWE map is available on a particular assimilation date at a resolution suitable for ingestion into the model.

1. Run the model from an initial state through the assimilation date and save the current model state on the assimilation date (maps of soil moisture, snow, interception, etc.).

2. Calculate ΔSWE map (measured – modeled) and replace modeled SWE with measured SWE.

3. Use machine learning (e.g., Random Forest) to reanalyze the snowpack energy balance current state based on the pre-assimilation snow state, so that properties of the post-assimilation snowpack like temperature, liquid water content, and albedo maintain their relationships to elevation, aspect, canopy cover, SWE depth, etc.

4. Infer and correct errors in antecedent water mass balance fluxes:

    4.1. Determine an "accounting period" over which the SWE error accumulated, beginning either with the most-recent snow assimilation date or the most-recent effectively snow-free period.

    4.2. Determine how much, if any, of the bulk SWE error can only be explained by precipitation data error:
        4.2.1. Reanalyze snow / rain partitioning on the accounting period by calculating how much more or less snow would have accumulated if the snow and rain temperature thresholds had been a few degrees warmer (if modeled SWE is too high) or cooler (if modeled SWE is too low).
        4.2.2. Add or subtract the modeled melt volume from the result of Step 4.2.1 to determine the largest SWE error that can plausibly be explained by errors in partitioning or ablation dynamics alone.
        4.2.3. If the result of Step 4.2.2 is greater in magnitude than the actual SWE error, then no precipitation error is inferred; otherwise, the unexplainable part of the SWE error is converted to a precipitation error (ΔP) by linear regression of SWE against cumulative precipitation on the accounting period.

    4.3. Determine a total error in terrestrial water input (TWI) by subtracting the total SWE error from the total precipitation error (ΔTWI = -ΔSWE when ΔP = 0).

    4.4. Partition ΔTWI between subsurface storage errors and antecedent runoff errors:
        4.4.1. Calculate a normalized runoff ratio between cumulative runoff (streamflow + channel storage) and net subsurface storage change (final soil moisture – initial soil moisture) on the accounting period.
        4.4.2. For each grid cell, calculate a ΔStorage error using the ΔSWE map and the runoff ratio from Step 4.4.1, and partition this error among grid cells along the downstream flow path in proportion to the current storage in each cell and inversely proportional to the upstream area above each cell.
        4.4.3. Similarly route ΔSWE errors through the stream network to derive a spatial pattern of channel storage errors in proportion to the pre-assimilation pattern of channel storage and upstream ΔSWE.
        4.4.4. Solve for the partitioning of total runoff error between antecedent streamflow and current channel storage so that the updated channel storage state is congruent with the latest streamflow time step.
        4.4.5. Determine a total streamflow error (ΔQ) by subtracting the channel storage error from the total runoff error calculated by application of the runoff ratio from Step 4.4.1 to ΔTWI.
        4.4.6. Reanalyze the antecedent streamflow hydrograph using weighted scaling factors so that (A.) the total change, ΔQ, agrees with Step 4.4.5, (B.) errors are roughly proportional to streamflow magnitude, (C.) the error is zero at the beginning of the accounting period, and (D.) the error scaling factors grow asymptotically according to $1 - e^{-\alpha * t}$ where α is a characteristic time derived from the rate of SWE change so that the timing of streamflow errors is consistent with the timing of snowpack changes.

5. Overwrite the antecedent streamflow timeseries and current snow and soil storage state files with the error-corrected versions; at this point, ΔSWE + ΔP + ΔQ + ΔStorage = 0, so mass is conserved.

6. Resume the model run on the assimilation date beginning from the error-corrected state, and go to Step 1 if additional (later) SWE data are available, or allow the model to finish running through the period of interest.

HISTORICAL DEMONSTRATION

The process-based assimilation method presented above is being robustly tested in a variety of scenarios. Real-time water supply forecast demonstrations of the DHSVM-WSF forecast platform are underway for the 2024 forecast season in two dozen watersheds across eight western states, with process-based assimilation of ASO lidar-based snow data in California, Colorado, and Wyoming. Backcast simulations using observed weather data are also useful for analyzing model performance under different conditions. This section presents a backcast demonstration of process-based assimilation using DHSVM-WSF and ASO SWE maps in the Tuolumne River watershed above Hetch Hetchy Reservoir in the Sierra mountains of California.

Model Setup and Calibration

DHSVM is customized and calibrated in each watershed using a combination of spatially explicit data sources and multi-objective Bayesian optimization. Please refer to the Mountain Hydrology white paper on DHSVM-WSF (Boardman 2023) for details on the datasets and techniques used to set up and calibrate the model. A few notes of particular relevance for the present discussion are summarized here. The model is currently being operated at 90 m resolution in the Tuolumne. Meteorological data are based on gridMET (Abatzoglou 2013), with precipitation redistributed using multiplier maps derived from ASO SWE patterns such that the total gridMET precipitation volume is preserved while the spatial distribution accounts for local wind and terrain effects. Historical daily inflow data for Hetch Hetchy Reservoir are smoothed and disseminated by Bruce McGurk (personal communication, February 2024). Historical ASO SWE data are archived on the ASO data portal (https://data.airbornesnowobservatories.com/). Key model parameters controlling subsurface transmissivity profiles, transpiration rates, and snow properties are calibrated using seven objective functions: daily Nash-Sutcliffe efficiency (NSE), log-scale NSE, yearly water yield percent error, April-July water yield percent error, monthly water yield root mean square error (RMSE), pixel-wise ASO SWE RMSE, and ASO SWE bulk volume percent error. Figure 4 shows a map of the land surface classification for DHSVM-WSF in the Hetch Hetchy watershed. In the snow zone, the dominant land cover types are alpine rock and conifer forest. The total watershed area above Hetch Hetchy Reservoir is 1,177 km^2, with elevations ranging from approximately 1,200 m to 4,000 m.

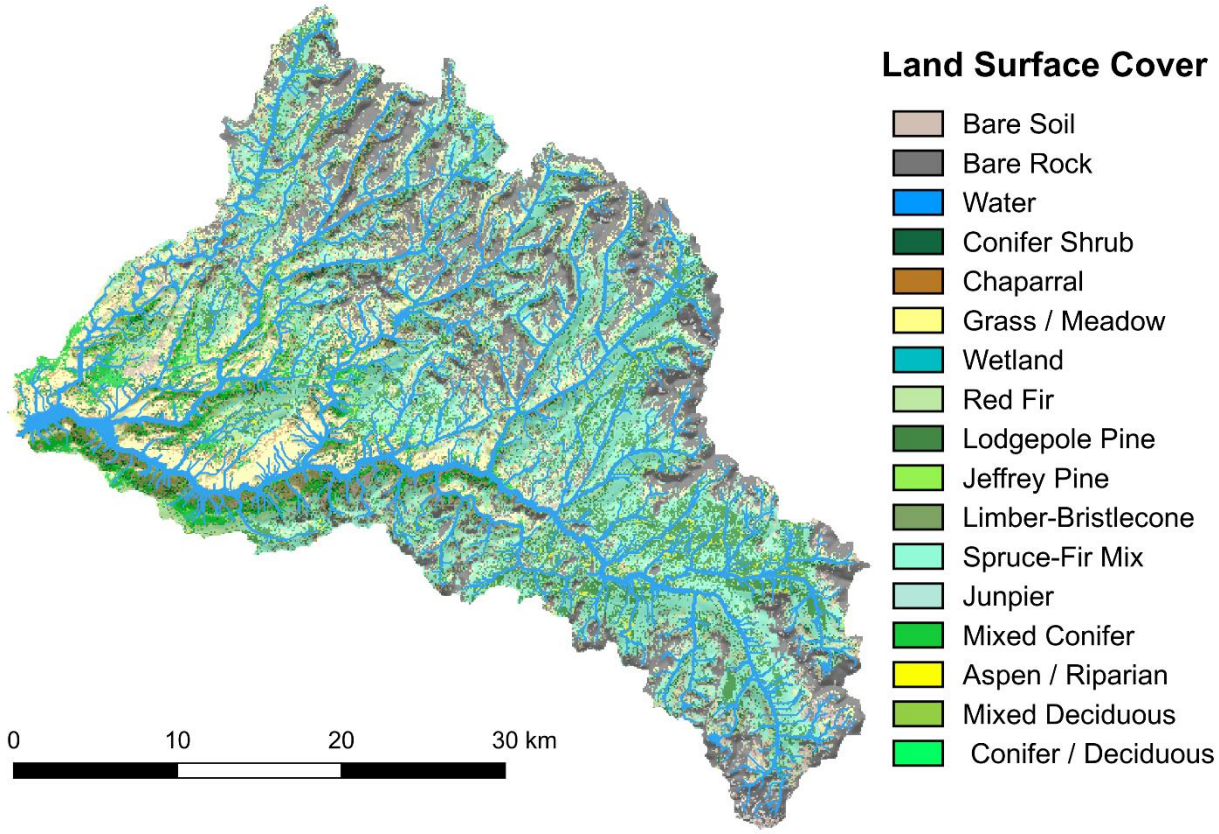

Figure 4. Map of the Tuolumne River watershed above Hetch Hetchy Reservoir, as modeled in DHSVM-WSF.

Backcast Simulations

The same DHSVM model is tested in three different modes to evaluate the historical performance of different snow data assimilation methods. The period of interest is defined as water years 2013 to 2023, encompassing all Tuolumne ASO flights up to the present year. There are a total of 72 ASO SWE maps available during this 11-year period, with a minimum of two SWE maps each year. A calibrated DHSVM model is initialized with two years of spin-up (water years 2011 and 2012), after which different assimilation methods are tested with the same set of SWE maps and other model inputs. In the "Open-Loop" mode, the model runs continuously without any snow data assimilation. In the "Direct Insertion" mode, the SWE map and other snowpack states are updated on each assimilation date (Steps 1-3 and 5-6 in the algorithm presented earlier), but no adjustments are made to non-snow components of the model. In the "Process-Based" assimilation mode, Step 4 of the assimilation algorithm is additionally implemented, which reanalyzes non-snow states and preserves the water mass balance. Figure 5 shows daily hydrographs resulting from application of all three model modes to the Hetch Hetchy watershed. Overall, the model captures observed variations in seasonal and year-to-year runoff magnitude and timing, though there are clear differences in the snowmelt response of the different assimilation modes.

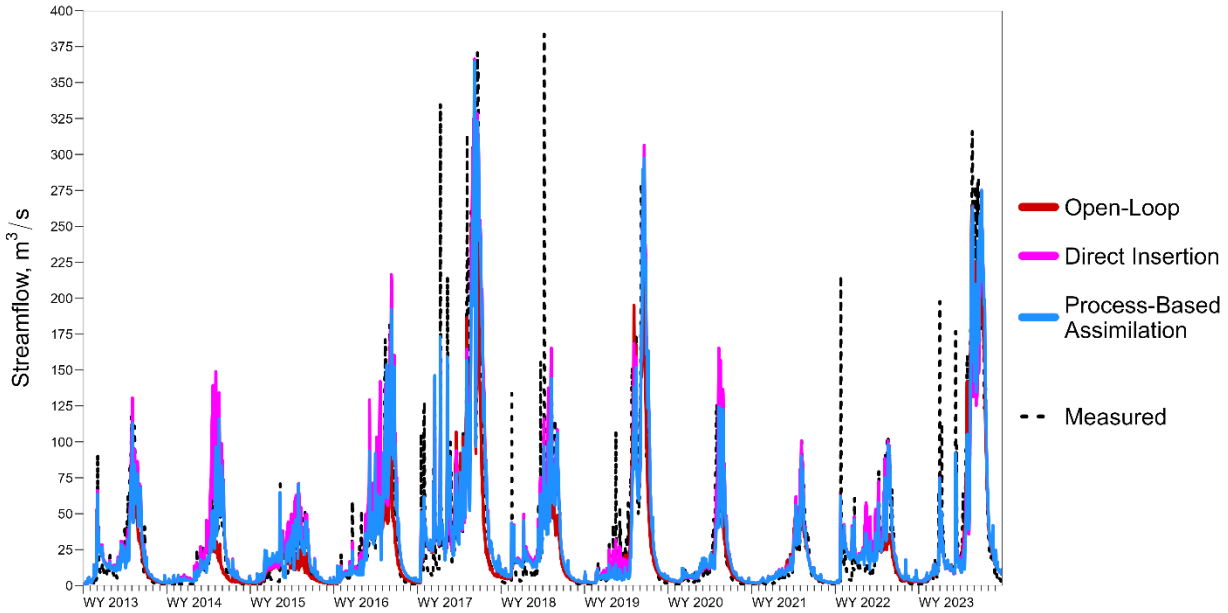

Figure 5. Daily measured and simulated hydrographs for Hetch Hetchy reservoir inflow, water years 2013-2023.

Process-Based Assimilation Results

The process-based snow data assimilation method consistently identifies and corrects non-snow errors in the hydrological model. Across the 72 assimilation dates (each representing an ASO SWE map), the mean basin SWE depth is 0.36 m, and the mean absolute model error is 0.022 m. The mean absolute percent error for the modeled SWE volume is 9% across all ASO acquisitions, with a range of -12% to +23%. The process-based assimilation method almost always preserves the water mass balance when increasing or decreasing the snowpack volume, with only 0.3% of total errors unaccounted. Out of the 72 SWE updates, all but 5 (93%) can be explained without invoking a precipitation data error, and 3 of the 5 inferred precipitation errors are in the extremely high-precipitation 2023 water year. Considering only the 47 assimilation dates with at least a 1 cm average SWE change (to avoid problems with exploding fractions), on average 85% of ΔTWI is partitioned to runoff errors, and 15% is partitioned to subsurface storage errors. Although pre-assimilation runoff does not directly affect post-assimilation runoff, accurately bias correcting the entire hydrograph is critical for predicting the water supply volume on a fixed time window. SWE assimilation frequently occurs during a period of interest (e.g., a May issue date for the April-July volume could assimilate a May SWE map), so it is critical to constrain the water mass balance as accurately as possible before and after each assimilation date. Moreover, even when the total subsurface storage volume does not change dramatically, the pattern may require significant updates. Figure 6 shows spatial changes resulting from just a small change to total volume (average +3 cm SWE, -1.5 cm soil moisture, remaining error attributed to runoff). In areas where the model melted the snow too quickly, some of the extra water was stored in the soil leading to erroneously high soil moisture, and vice versa, which is corrected by the process-based assimilation method.

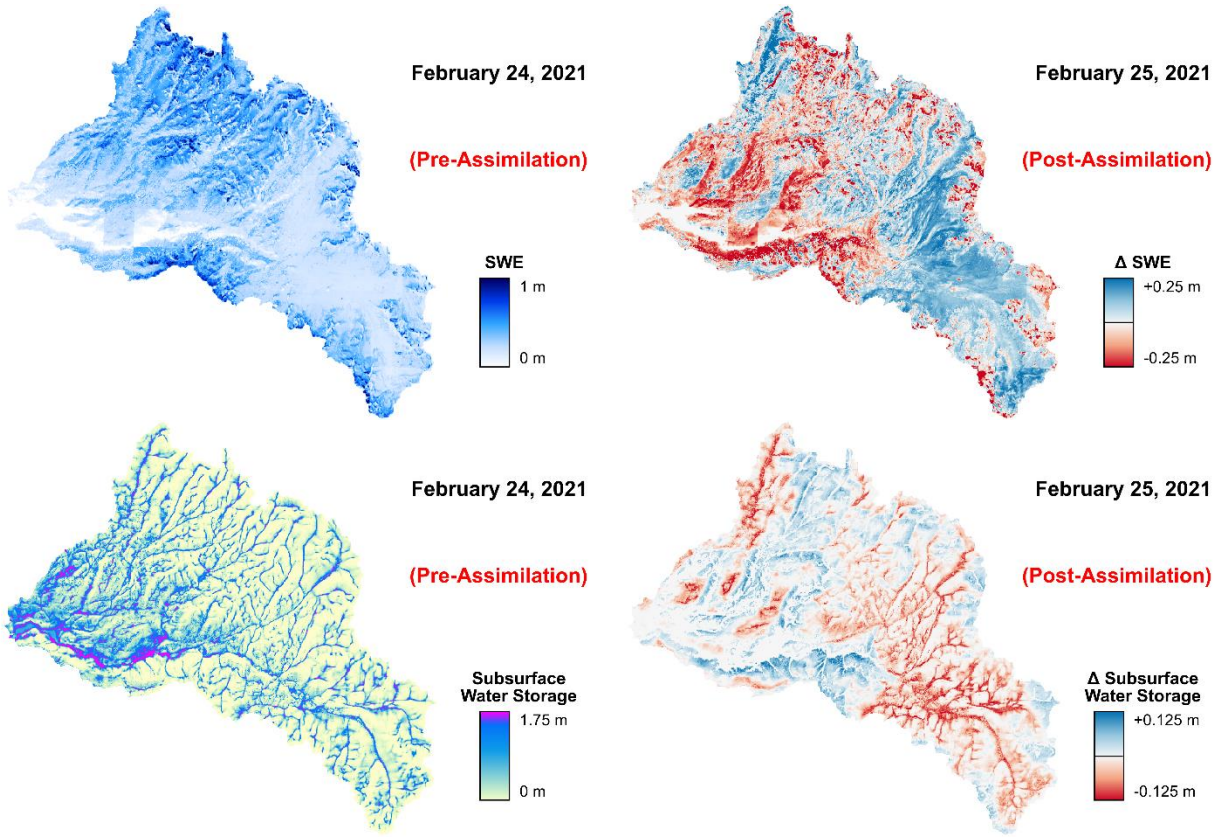

Figure 6. Example of pre- and post-assimilation snow and subsurface water storage, showing changes in subsurface water storage to correct inferred errors resulting from imperfect simulation of the snowpack dynamics. Note how propagating spatial SWE errors through the subsurface results in complementary reverse error patterns, except that subsurface errors are concentrated in topographic convergence zones due to antecedent groundwater flow.

Comparison of Assimilation Methods

Process-based assimilation outperforms both direct insertion and the open-loop model by a wide margin (Table 1). All three simulation modes use the same baseline calibrated DHSVM model, and this model consistently achieves a satisfactory daily NSE of 0.82-0.84 regardless of the assimilation method. The open-loop model (no assimilation) has a reasonable mean absolute percent error (MAPE) of 11% for the yearly runoff volume, suggesting that the gridMET precipitation volume is reasonable, but sometimes the model simulates snowmelt runoff too early and sometimes too late, causing a high MAPE of 27% for the target April-July volume. Direct insertion of the ASO SWE maps slightly improves the April-July volume error, but at the cost of drastically higher (29%) yearly MAPE. This is an unambiguous result of "double counting" water in some years (the same snow is melted twice) and "black holes" in other years (extra snow is subtracted instead of added to streamflow). The process-based assimilation method produces by far the lowest April-July error (7% MAPE), with only marginally higher yearly error compared to the open-loop model and the highest daily NSE. For the April-July volume, process-based assimilation can reduce backcast error by 74% relative to the open-loop model and by 63% relative to direct insertion of the same SWE data.

Assimilation Method	Daily NSE	Yearly Mean Absolute % Error	April – July Mean Absolute % Error
Open-Loop (None)	0.83	11%	27%
Direct Insertion	0.82	29%	19%
Process-Based	0.84	14%	7%

Table 1. Summary of backcast error metrics for different snow data assimilation methods.

Year-to-Year Variability

Process-based assimilation is much more resilient to climatic variability compared to direct insertion. In Table 2, the backcast April-July volume error is always the highest in the open-loop model, which demonstrates the consistent value of incorporating measurement-based SWE maps. The direct insertion method produces the smallest error in 4 out of 11 years, and process-based assimilation produces the smallest error in the other 7 years. However, when direct insertion fails, it fails catastrophically. Direct insertion of SWE maps in 2014, 2015, and 2020 produces volumetric runoff errors of 70%, 40%, and 30% respectively, an average of 35% more error than the process-based method in those same years. In the years that direct insertion outperforms process-based assimilation, the latter only has 6% more error on average. It is worth noting that the years when direct insertion causes the largest errors are also some of the driest years. Direct insertion overestimates the water supply in dry years by as much as 70%, which could have severe negative impacts on water management and drought preparedness. Overall, process-based snow data assimilation outperforms the other simulation modes both on average and during hydrologic extremes.

Water Year	**Number of Snow Assimilation Dates**	**AJ Volume** *km^3 (TAF)*	**AJ % Error** *Open-Loop*	**AJ % Error** *Direct Insertion*	**AJ % Error** *Process-Based*
2013	6	0.42 (343)	-24%	+5%	-8%
2014	11	0.33 (268)	-51%	+70%	+16%
2015	10	0.23 (187)	-50%	+40%	+12%
2016	13	0.71 (579)	-29%	+20%	+3%
2017	9	1.57 (1,274)	-15%	+6%	0%
2018	2	0.62 (505)	-30%	-1%	-13%
2019	5	1.13 (916)	-19%	-1%	-6%
2020	3	0.37 (303)	-11%	+30%	+5%
2021	2	0.29 (236)	-7%	+10%	0%
2022	5	0.38 (308)	-44%	+4%	-7%
2023	6	1.58 (1,277)	-14%	-16%	-11%
Mean Absolute:	**72**	**0.69 (563)**	**27%**	**19%**	**7%**

Table 2. Backcast April-July volume percent error for different snow data assimilation methods.

CONCLUSION

Process-based assimilation merges the salient aspects of remotely sensed snow data and distributed physical hydrological models. Conventional (e.g., ensemble-based) assimilation methods are unsuitable for remote sensing snow data, because errors are spatial and not just volumetric. Direct insertion of remotely sensed snow data succeeds at matching observed spatial distributions, but it also violates the most fundamental principle of physical hydrological modeling: the water mass balance. In process-based assimilation, an approximate time-reversible model can be constructed from relationships between simulated pre-assimilation water mass balance fluxes. A number of other approaches might be conceived to partition SWE errors between model components, but since the model itself is designed specifically to partition the water mass balance, it is intuitively appealing to use the model's own outputs as a first-order approximation to its internal errors. Furthermore, process-based assimilation method can be implemented during calibration using historical data, so the validity of the error partitioning method is implicitly tested and refined. In this sense, the whole assimilation system becomes a single internally consistent structure encompassing the modeled processes and their responses to observed data. Results presented here demonstrate the large potential for improvement from process-based SWE map assimilation, including reduced error and improved reliability during drought. With growing emphasis on hydrological models that can learn from data while retaining physical interpretability, process-based assimilation offers a promising path forward.

ACKNOWLEDGEMENTS

I would like to thank Bruce McGurk for curating the Hetch Hetchy Reservoir inflow data, the ASO team for making the ASO SWE maps publicly available, and John Abatzoglou for making the gridMET meteorological data publicly available. I am grateful to Zhuoran Duan, Mark Wigmosta, and Adrian Harpold for their guidance on hydrologic modeling in general and DHSVM specifically. I am grateful to the Hetch Hetchy Reservoir team for valuable insight into water supply forecasting and Tuolumne River hydrology. Data and models supporting the findings presented in this paper are available from the author upon reasonable request.

REFERENCES

Abatzoglou, J. T. (2013). Development of gridded surface meteorological data for ecological applications and modelling. *International Journal of Climatology*, 33(1), 121–131. https://doi.org/10.1002/joc.3413

Boardman, E. N. (2023). DHSVM-WSF: The Distributed Hydrology Soil Vegetation Model for Water Supply Forecasts. Mountain Hydrology LLC, White Paper 001v3. https://mountainhydrology.com/mh001_dhsvm-wsf/

Church, J. E. (1933). Snow Surveying: Its Principles and Possibilities. *Geographical Review*, 23(4), 529–563. https://doi.org/10.2307/209242

Gochis, D. J., Barlage, M., Cabell, R., Casali, M., Dugger, A., FitzGerald, K., McAllister, M., McCreight, J., RafieeiNasab, A., Read, L., Sampson, K., Yates, D., & Zhang, Y. (2020). The WRF-Hydro® modeling system technical description, Version 5.2.0. NCAR Technical Note. https://ral.ucar.edu/sites/default/files/public/projects/wrf-hydro/technical-description-user-guide/wrf-hydrov5.2technicaldescription.pdf

Hedrick, A. R., Marks, D., Havens, S., Robertson, M., Johnson, M., Sandusky, M., Marshall, H.-P., Kormos, P. R., Bormann, K. J., & Painter, T. H. (2018). Direct Insertion of NASA Airborne Snow Observatory-Derived Snow Depth Time Series Into the iSnobal Energy Balance Snow Model. *Water Resources Research*, 54(10), 8045–8063. https://doi.org/10.1029/2018WR023190

Nolin, A. W. (2010). Recent advances in remote sensing of seasonal snow. *Journal of Glaciology*, 56(200), 1141–1150. https://doi.org/10.3189/002214311796406077

Pagano, T., Garen, D., & Sorooshian, S. (2004). Evaluation of Official Western U.S. Seasonal Water Supply Outlooks, 1922–2002. *Journal of Hydrometeorology*, 5(5), 896–909. https://doi.org/10.1175/1525-7541(2004)005<0896:EOOWUS>2.0.CO;2

Painter, T. H., Berisford, D. F., Boardman, J. W., Bormann, K. J., Deems, J. S., Gehrke, F., Hedrick, A., Joyce, M., Laidlaw, R., Marks, D., Mattmann, C., McGurk, B., Ramirez, P., Richardson, M., Skiles, S. M., Seidel, F. C., & Winstral, A. (2016). The Airborne Snow Observatory: Fusion of scanning lidar, imaging spectrometer, and physically-based modeling for mapping snow water equivalent and snow albedo. *Remote Sensing of Environment*, 184, 139–152. https://doi.org/10.1016/j.rse.2016.06.018

Schaake, J. C. Jr., & Peck, E. L. (1985). Analysis of Water Supply Forecast Accuracy. Proceedings of the 53rd Annual Western Snow Conference. 53rd Annual Western Snow Conference. sites/westernsnowconference.org/PDFs/1985Schaake.pdf

Stillinger, T., Costello, C., Bales, R. C., & Dozier, J. (2021). Reservoir Operators React to Uncertainty in Snowmelt Runoff Forecasts. *Journal of Water Resources Planning and Management*, 147(10), 06021010. https://doi.org/10.1061/(ASCE)WR.1943-5452.0001437

Wigmosta, M. S., Vail, L. W., & Lettenmaier, D. P. (1994). A distributed hydrology-vegetation model for complex terrain. *Water Resources Research*, 30(6), 1665–1679. https://doi.org/10.1029/94WR00436