Supporting Information

A near-term iterative forecasting system successfully predicts reservoir hydrodynamics and partitions uncertainty

R. Quinn Thomas, Renato J. Figueiredo, Vahid Daneshmand, Bethany J. Bookout, Laura K. Puckett, Cayelan C. Carey

Pages: 19

Text: 3

Tables: 3

Figures: 4

Supporting Information A. Detailed description of the data assimilation methods used in FLARE.

The data assimilation in FLARE used the ensemble Kalman Filter (EnKF) with state augmentation to calibrate parameters (following the methods of [*Zhang et al. 2017*]). The EnKF state matrix had *M* ensemble members, each with *K* model depths (state variables) and *P* number of parameters (an augmentation of the states by including parameters), resulting in an $M \times (K + P)$ matrix.

The EnKF was initialized with a set of *M* ensemble members, in which each ensemble member *i* had a vector of modeled water temperatures at *K* depths at the *0*th time (x_0^i) and a vector of *P* parameters (α_0^i). For the first day of data assimilation at the beginning of the spin-up period only, the values of x_0^i were initialized with observed sensor temperatures and linear interpolation was used to initialize the modeled depths that did not have observations.

For this application, *P* was three because three General Lake Model (GLM) parameters were calibrated for Falling Creek Reservoir: SW_factor, LW_factor, zone1temp, and zone2temp. These parameters were chosen based on a one-step-at-a-time (OAT) global sensitivity analysis of all GLM parameters [*Morris 1991*]. α_0^i was initialized using a random draw for each ensemble from a parameter-specific uniform distribution. For every sequential day in the spin-up forecasting and forecasting periods, a new vector of parameters for each ensemble member (α_t^{i-}) was created by adding a normal random variable centered at 0 with a specified covariance (ϕ) to the previous day's parameter values (eqn. SI.1). The negative sign in the α_t^{i-} signifies a parameter vector before updating using assimilated observations, following eqn. SI.1:

$$\alpha_t^{i-} = \alpha_{t-1}^{i} + MVN(0, \Theta) \qquad \text{(eqn. SI.1)}$$

The covariance (Θ) was constant throughout assimilation and was set to be small but non-zero to allow the *P* parameters to adjust over time and improve the model calibration.

Every day, the observed meteorology from the previous 24 hours was pulled from the GitHub repository and processed to generate a matrix of hourly meteorological inputs for GLM. This matrix was combined with the other model driver data (inflow rate, mean historical 5-year inflow water temperature, and mean historical 5-year outflow rate) to create a driver matrix (D_t^i) for each ensemble member. The GLM inputs did not differ among the ensembles when assimilating observations using historical observations.

The vector of modeled water temperature for each depth from the previous day (x_{t-1}^i), the parameter vector (α_t^{i-}), and the last 24 hours of driver data (D_t^i) were used to initialize and run a 1-day simulation of the GLM for each ensemble member, $G\left(x_{t-1}^i, \alpha_t^{i-}, D_t^i\right)$. Process uncertainty was added to the water temperature predictions from the GLM following eqn. SI.2 to create predictions of water temperature with process uncertainty for each depth:

$$x_t^{i-} = G\left(x_{t-1}^i, \alpha_t^{i-}, D_t^i\right) + MVN(0, \Sigma_t) \text{ (eqn. SI.2)}$$

where x_t^{i-} is the $K \times 1$ vector of predicted water temperatures at the modeled depths for the *i*th ensemble member at time *t*. $\mathrm{MV}N(0, \Sigma_t)$ is a random draw from a multivariate normal distribution with a mean of 0 and the covariance matrix at time *t* (Σ_t).

The Σ_t matrix evolved through data assimilation, as the model predictions prior to updating (x_t^{i-}) improved or degraded over time. This allowed for the process uncertainty to reflect the performance of model predictions over a specified time period (a 30-day window in our application for Falling Creek Reservoir). The first 30 days of assimilation used to generate the Σ_t matrix so that the Σ_t during that period did not evolve and was a diagonal matrix with a constant variance for all depths (0.5 °C). After the first 30 days, a 30-day running covariance matrix at the observed depths (Σ_t^*) was calculated as the residual of the predictions prior to updating, following eqn. SI.3:

$$\boldsymbol{\Sigma}_{\boldsymbol{t}}^* = \frac{1}{V}\sum_{l=t}^{V}(\overline{x_{t-l}} - y_{t-l})(\overline{x_{t-l}} - y_{t-l}) \quad \text{(eqn. SI.3)}$$

$\boldsymbol{\Sigma}_{\boldsymbol{t}}^*$ was used to calculate $\boldsymbol{\Sigma}_{\boldsymbol{t}}$ by linearly interpolating the variances and covariances between depths in $\boldsymbol{\Sigma}_{\boldsymbol{t}}^*$. In eqn. SI.3, V is the number of previous days included in the covariance matrix (here, 30).

If data were not available to update the model states due to missing sensor data, the states were not updated and $x_t^i = x_t^{i-}$. Otherwise, we calculated the covariance among states in the ensemble members ($\boldsymbol{C}_{\boldsymbol{xx}}$) using eqn. SI.4:

$$\boldsymbol{C}_{\boldsymbol{xx}} = \frac{1}{M-1}\sum_{i=1}^{M}(x_t^{i-} - \overline{x_t})(x_t^{i-} - \overline{x_t}) \quad \text{(eqn. SI.4)}$$

where $\overline{x_t}$ was the mean temperature at each modeled depth across ensemble members. The $\boldsymbol{C}_{\boldsymbol{xx}}$ matrix represents the estimated model error. Similarly, we calculated the covariance among parameters and states in the ensemble members ($C_{\alpha x}$) to estimate the relationship between parameters and model predictions using eqn. SI.5:

$$\boldsymbol{C}_{\boldsymbol{\alpha x}} = \frac{1}{M-1}\sum_{i=1}^{M}(\alpha_t^{i-} - \overline{\alpha_t})(x_t^{i-} - \overline{x_t}) \quad \text{(eqn. SI.5)}$$

where $\overline{\alpha_t}$ was the mean across ensemble members for each parameter in the parameter vector.

Next, to quantify uncertainty in the observations, we added normally-distributed noise to the vector of observations at time t (y_t) using the observation covariance matrix ($\boldsymbol{R}$) (eqn. SI.6):

$$\hat{y}_t^i = y_t + MVN(0, \boldsymbol{R}) \quad \text{(eqn. SI.6)}$$

where $\hat{y}_t^i$ is the vector of observations with uncertainty added. In our application, the observational uncertainty was equal for all depths and not correlated among depths, and thus the $\boldsymbol{R}$ matrix was diagonal.

The model states (water temperatures at specific depths) and parameter updating using the observations first required calculating the Kalman gain for the states ($\boldsymbol{K_x}$) and parameters ($\boldsymbol{K_\alpha}$) following eqn. SI.7:

$$\begin{bmatrix}\boldsymbol{K_x}\\ \boldsymbol{K_\alpha}\end{bmatrix} = \begin{bmatrix}\boldsymbol{C_{xx}H^T}(\boldsymbol{HC_{xx}H^T} + \boldsymbol{R})^{-1}\\ \boldsymbol{C_{\alpha x}H^T}(\boldsymbol{HC_{xx}H^T} + \boldsymbol{R})^{-1}\end{bmatrix} \text{ (eqn. SI.7)}$$

where $\boldsymbol{H}$ is a matrix in which each row corresponds to a depth with an observation and each column represents each of the modeled depths. The column that matched the depth of the particular row's observation had a value of 1 while all other columns had a value of 0. Each row only had a single 1. T represents the transpose of the $\boldsymbol{H}$ matrix.

The Kalman gain represented the proportional adjustment of the GLM model output based on the difference between the model predictions of water temperature and the sensor observations. A Kalman gain value of 1 is associated with a full adjustment of the model state to match an observation (likely due to low or near-zero observational uncertainty; $\boldsymbol{R}$), while a value of zero has no adjustment of the modeled state. The full matrix of the Kalman gain included the direct updating of water temperature at a particular depth based on the comparison to sensor observations at that depth and on the covariance of model states across depths (i.e., a large update in one depth influenced the update of another depth if there was high correlation between those two specific depths). This allowed the Kalman gain to update depths without sensor observations because they were correlated with observed depths in the model predictions in $\boldsymbol{C_{xx}}$.

Finally, the corrupted states and parameters were updated by adding the state gain $\boldsymbol{K_x}(\hat{y}_t^i - \boldsymbol{H}x_t^{i-})$ and parameter gain $\boldsymbol{K_\alpha}(\hat{y}_t^i - \boldsymbol{H}x_t^{i-})$, using eqn. SI.8:

$$\begin{bmatrix}x_t^i\\ \alpha_t^i\end{bmatrix} = \begin{bmatrix}x_t^{i-}\\ \alpha_t^{i-}\end{bmatrix} + \begin{bmatrix}\boldsymbol{K_x}(\hat{y}_t^i - \boldsymbol{H}x_t^{i-})\\ \boldsymbol{K_\alpha}(\hat{y}_t^i - \boldsymbol{H}x_t^{i-})\end{bmatrix} \text{ (eqn. SI.8)}$$

The Kalman gain thus updated the model states for which there were corresponding observations and updated the model states that did not have corresponding observations based on the correlation between the observed and unobserved states. Similarly, the parameters were updated based on their correlation with the observed states.

References

Morris, M.D. (1991), Factorial Sampling Plans for Preliminary Computational Experiments. Technometrics 33(2), 161-174, https://doi.org/10.1080/00401706.1991.10484804

Zhang, H., Hendricks Franssen, H.-J., Han, X., Vrugt, J.A. and Vereecken, H. (2017), State and parameter estimation of two land surface models using the ensemble Kalman filter and the particle filter. Hydrol. Earth Syst. Sci. 21(9), 4927-4958, https://doi.org/10.5194/hess-21-4927-2017

Supporting Information B. Description of how the NOAA GEFS forecasts were spatially and temporally-downscaled.

The overarching goal of the spatial and temporal downscaling was to adjust the 1 x 1° spatial resolution and 6-hour temporal resolution NOAA GEFS forecasts to represent the reservoir's local meteorological conditions at a 1-hour temporal resolution.

First, we used historical GEFS forecasts and 1-minute scale observational data measured at the reservoir from 6 April – 6 December 2018 as the "training data" for the spatial downscaling [*Carey et al. 2019*]. We aggregated both the NOAA GEFS and the observed meteorology to the daily scale by averaging all observations (except for precipitation, which was summed) and matched the data by date. In this training dataset, we only used the first day of each historical 16-day NOAA GEFS forecast because it contained the lowest spread among NOAA GEFS ensemble members and was mostly likely to represent any consistent offsets between the 1×1° forecast and the local conditions.

To spatially-downscale temperature, relative humidity, wind speed, shortwave radiation, and longwave radiation, we estimated the linear relationship between the daily observation and forecast data in the training dataset (Supporting Information Table 2). We then applied this linear model to each day of the 16-day forecast. We set downscaled values for each variable that was less than zero to zero and values of relative humidity greater than 100 to 100. This resulted in a spatially-downscaled NOAA GEFS forecast product at the daily time scale.

To temporally-downscale the spatially-downscaled temperature, relative humidity, and wind speed forecasts from the daily to 1-hour resolution, we first used the difference between the pre-spatially downscaled NOAA GEFS 6-hour forecast and its daily mean to convert the daily spatially-downscaled forecast to its original 6-hour resolution. We used a monotone Hermite spline method to obtain hourly values from the 6-hour values. Before applying the spline method within the first 6-hour period, we used the observed meteorology as the 0-hour variable and the downscaled forecast as the 6-hour value. This allowed for a smooth transition between the observed meteorology used in data assimilation and the downscaled forecast.

To temporally-downscale shortwave radiation from the spatially-downscaled daily resolution to 1-hour resolution, we calculated the potential top-of-atmosphere solar radiation for each hour to determine a scaling factor between hourly shortwave radiation and the mean daily potential shortwave radiation [following the solar_geom.R function in *Dietze 2017*]. We used this ratio to convert the daily downscaled shortwave radiation to the 1-hour resolution.

To temporally-downscale longwave radiation from the spatially-downscaled daily resolution to 1-hour resolution, we first used the relative difference between the pre-spatially downscaled NOAA GEFS 6-hour forecast and its daily mean to convert the daily spatially-downscaled forecast to its original 6-hour resolution. We then applied the 6-hour mean value to each hour within that time window.

Precipitation was only spatially-downscaled. We first calculated the ratio of the forecasted precipitation to observed precipitation in the training data. Then, we multiplied each NOAA GEFS 6-hourly forecasts of precipitation by this ratio.

Finally, we represented uncertainty in the spatial and temporal-downscaling process by adding random noise to each downscaled 1-hour forecast. To add the random noise, we first applied the spatial and temporal downscaling process described above to the NOAA GEFS forecast used in the training data. Second, we calculated the residuals between the observed meteorology and the downscaled NOAA GEFS forecast at the 1-hour resolution for temperature, relative humidity, wind speed, shortwave radiation, and longwave radiation. This resulted in a set of residuals for each variable (except precipitation) within each hour. Third, we used the residuals to determine the covariance of residuals among variables across all hours in the training dataset (Supporting Information Table 3). Finally, to add noise to each hour of a 16-day forecast, we used this covariance to draw values for each variable from a multivariate normal distribution that was centered at the downscaled values. By using the multivariate normal distribution, the added noise reflects the downscaling uncertainty that is not independent among variables. In total, we generated 21 random draws from the downscaling uncertainty for each of the 21 downscaled NOAA GEFS ensembles.

References

Carey, C.C., Bookout, B.J., Lofton, M.E. and McClure, R.P. (2019), Time series of high-frequency meteorological data at Falling Creek Reservoir, Virginia, USA 2015-2018., Environmental Data Initiative, https://doi.org/10.6073/pasta/68de79f732a9f3d2a686dda2eeb8197d

Dietze, M.C. (2017), Prediction in ecology: a first-principles framework. Ecol Appl 112(1), 6252-6213, https://doi.org/10.1002/eap.1589

Supporting Information C: A description of the sensor array at the reservoir and wireless data transmission methods.

We measured the water temperature profile in Falling Creek Reservoir on 1-m intervals from the surface (0.1 m depth) to just above the sediments at 9 m at the deepest site of the reservoir with NexSens T-Node FR thermistors (NexSens Technology, Inc.; Fairborn, Ohio, USA; [*Carey et al. 2019b*]. Thus, we had sensor observations for 0.1 m, 1 m, 2 m, 3 m, 4 m, 5 m, 6 m, 7 m, 8 m, and 9 m. The thermistor string was factory-calibrated and verified against a NIST-traceable thermistor to meet measurement accuracy of ±0.075^{o}C. A Campbell Scientific (Logan, Utah, USA) research-grade meteorological station deployed on the dam of the reservoir measured shortwave radiation, longwave radiation, air temperature, relative humidity, rainfall, wind speed, and barometric pressure [*Carey et al. 2019a*]. These meteorological variables were measured every minute and then downsampled (temperature, wind speed, humidity), averaged (shortwave and longwave), or summed (precipitation) to the hourly scale to serve as driver data for the GLM model (Supporting Information Table 1).

The water temperature and meteorological sensor data were staged on Campbell Scientific data loggers on-site as measurements were retrieved, and transmitted daily to cloud storage. The sensor gateway attached to the Campbell Scientific data loggers ran the Ubuntu Linux software distribution, as well as software applications and scripts that were developed to perform data transfer and management functions including: 1) retrieve data from the logger using Campbell Scientific interfaces, 2) check cellular modem connectivity and reset modules as needed; and 3) reliably upload sensor data updates to appropriate repositories on cloud storage using the git client. Data were structured as a time series, with measurements appended as lines to a comma-separated values (CSV) file. Data transfers used Git (https://git-scm.com), an open-source distributed version control system, for efficient and reliable updates with minimum bandwidth usage, such that only the data collected since the last successful transfer were sent from the gateway to the cloud server. The gateway also ran a virtual private network (VPN) open-source software, IPOP (IP-over-P2P) to provide authentication and encryption [*Ganguly et al. 2006*], thereby providing a secure data transfer.

We measured the inflow discharge rate of the primary tributary entering into FCR through a weir with an INW Aquistar PT2X pressure sensor (INW, Kirkland, Washington, USA), which recorded the water temperature and water level [*Carey et al. 2018*]. We used the water level to calculate the mean daily discharge rate following [*Gerling et al. 2014*] and set the outflow discharge rate to the inflow discharge rate as the reservoir was maintained at a constant water level through the study.. Because we were unable to wirelessly connect the weir sensor to the cloud to transmit the inflow discharge data in real-time, we averaged the previous five years' data measured on a given day to serve as driver data for forecasting.

References

Carey, C.C., Bookout, B.J., Lofton, M.E. and McClure, R.P. (2019a), Time series of high-frequency meteorological data at Falling Creek Reservoir, Virginia, USA 2015-2018., Environmental Data Initiative, https://doi.org/10.6073/pasta/68de79f732a9f3d2a686dda2eeb8197d

Carey, C.C., Bookout, B.J. and Woelmer, W.M. (2019b), Time series of high-frequency sensor data measuring water temperature, dissolved oxygen, conductivity, specific conductivity, total dissolved solids, chlorophyll a, phycocyanin, and fluorescent dissolved organic matter at discrete depths in Falling Creek Reservoir, Virginia, USA in 2018, Environmental Data Initiative, https://doi.org/10.6073/pasta/128777b056a4b2bfba99ff8e780e2d59

Carey, C.C., Gerling, A.B., McClure, R.P., Lofton, M.E. and Bookout, B.J. (2018), Discharge time series for the primary inflow tributary entering Falling Creek Reservoir, Vinton, Virginia, USA 2013-2018., Environmental Data Initiative, https://doi.org/10.6073/pasta/64ff214b987da2997f5c823b156b3334

Ganguly, A., Agrawal, A., Boykin, P.O. and Figueiredo, R. (Year), IP over P2P: enabling self-configuring virtual IP networks for grid computing. Proceedings 20th IEEE International Parallel & Distributed Processing Symposium. 10.1109/IPDPS.2006.1639287

Gerling, A.B., Browne, R.G., Gantzer, P.A., Mobley, M.H., Little, J.C. and Carey, C.C. (2014), First report of the successful operation of a side stream supersaturation hypolimnetic oxygenation system in a eutrophic, shallow reservoir. Water Res 67, 129-143, https://doi.org/10.1016/j.watres.2014.09.002

Supporting Information Table 1. Meteorological sensors deployed on the dam at the Falling Creek Reservoir as part of a research-grade Campbell Scientific weather station that collected driver data for the General Lake Model.

Sensors deployed at the reservoir	Meteorological variables measured	Measurement precision
Rotronic Hydroclip2 HC2S3-L Temperature and Relative Humidity Probe with RM Young 10 plate Solar Radiation Shield	Air Temperature at 2 m	-50 - 100℃ ± 0.1
	Relative Humidity at 2 m	0 - 100% ± 1.3
RM Young 05103-L Wind Monitor	Wind Speed at 4 m	0 - 100 m/s ± 0.3
Hukseflux NR01 4-component Net Radiometer	Surface Downward Shortwave Radiation Flux	0 - 2000 W/m^2 ± 10%
	Surface Downward Longwave Radiation Flux	0 - 1000 W/m^2 ± 10%

Supporting Information Table 2. The slope, intercept, and R^2 for the relationship between the first day of each NOAA GEFS forecast for the grid cell that contains Falling Creek Reservoir and the observed meteorology from the on-site weather station, as described in Supporting Information B.

	Slope	Intercept	R^2
Air temperature	0.97	10.3	0.95
Relative humidity	1.0	-1.4	0.55
Wind speed	0.53	0.68	0.46
Shortwave radiation	0.77	7.40	0.81
Longwave radiation	0.96	43.5	0.94

Supporting Information Table 3. Covariance matrix describing the relationships among residuals from the observed meteorology and downscaled NOAA GEFS forecasts (see Supporting Information B).

	Air temperature	Wind speed	Relative humidity	Shortwave radiation	Longwave radiation	Rain
Air temperature	2.26	0.05	-5.12	17.64	-0.11	0
Wind speed	0.05	0.26	-0.54	3.45	-1.98	0
Relative humidity	-5.12	-0.54	80.26	-75.29	16.29	0
Shortwave radiation	17.64	3.45	-75.29	1361.29	-231.28	0
Longwave radiation	-0.11	-1.98	16.29	-231.28	147.29	0
Rain	0	0	0	0	0	0

Supporting Information Figure 1. Values for the three calibrated parameters: a) shortwave factor, b) mean zone 1 sediment temperature, and c) mean zone 2 sediment temperature from each ensemble member during the combined spin-up and forecasting periods of the study.

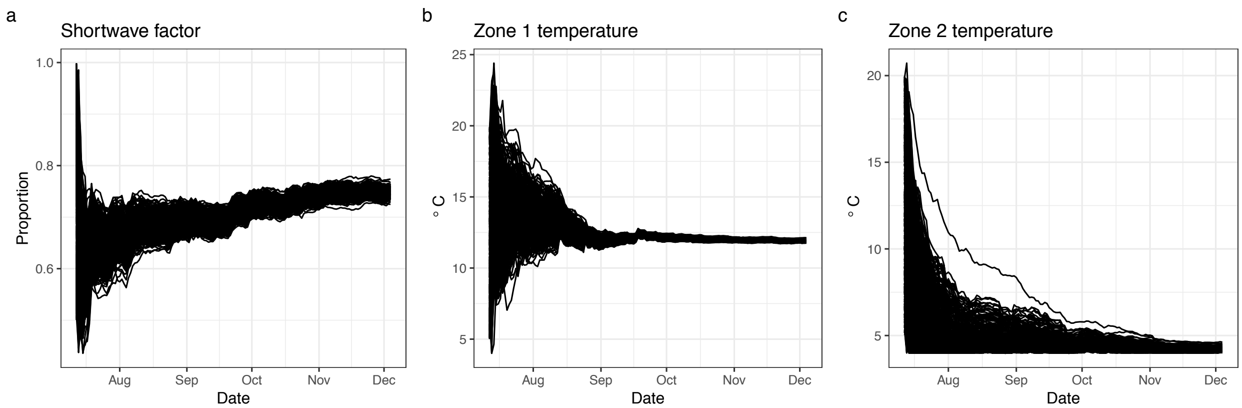

Supporting Information Figure 2. The relative contributions of the individual sources of uncertainty (left axis) to the total forecast uncertainty (right axis, orange line) varies through the 16-day forecast horizon. This forecast was initialized on 1 September 2018 and is one of three 16-day forecasts (with Supporting Information Figure 3 and Supporting Information Figure 4) that were averaged to create Figure 6. Two depths are shown (0.1 m – a, b; 8.0 m – c, d) and the relative contributions of initial condition uncertainty without (left) and with (right) gaps in water temperature sensor observations are shown in the two columns.

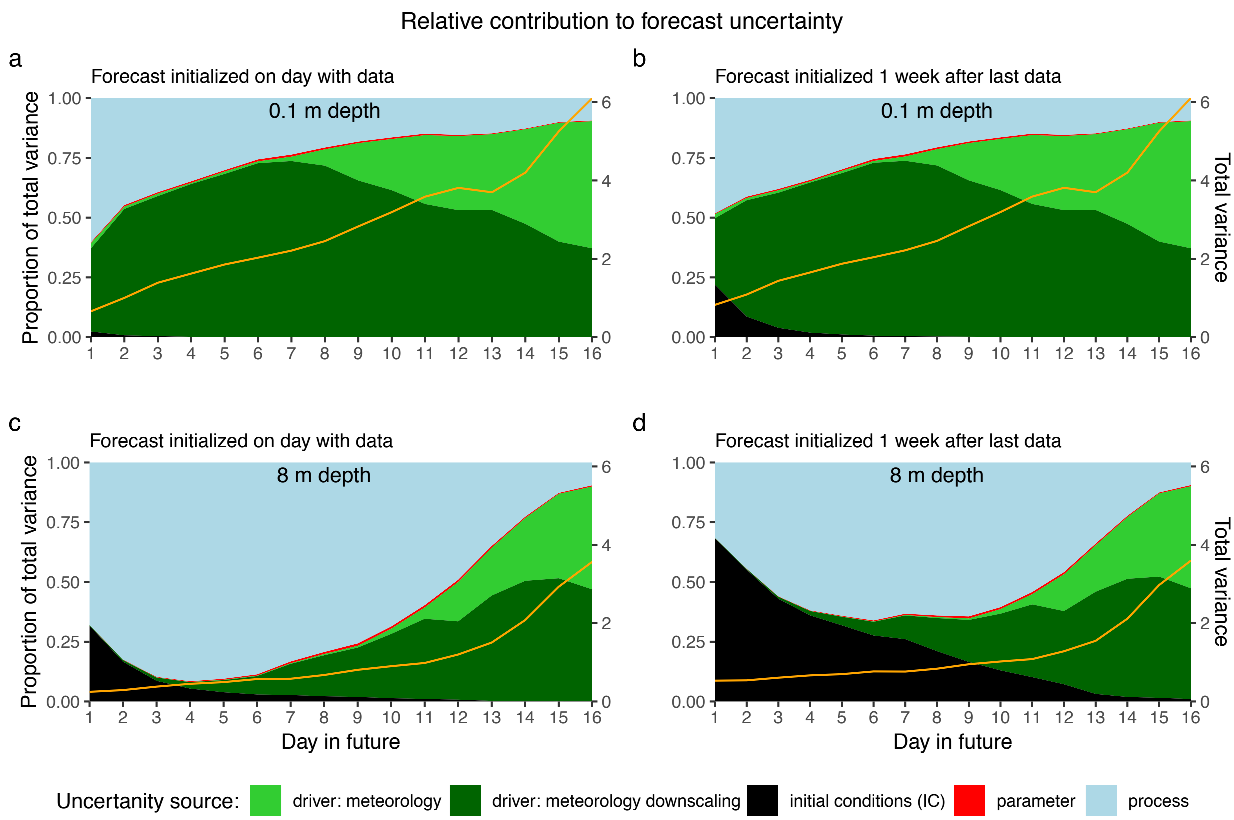

Supporting Information Figure 3. The relative contributions of the individual sources of uncertainty (left axis) to the total forecast uncertainty (right axis, orange line) varies through the 16-day forecast horizon. This forecast was initialized on 18 October 2018, three days prior to turnover, and is one of three 16-day forecasts (with Supporting Information Figure 2 and Supporting Information Figure 4) that were averaged to create Figure 6. Two depths are shown (0.1 m – a, b; 8.0 m – c, d) and the relative contributions of initial condition uncertainty without (left) and with (right) gaps in water temperature sensor observations are shown in the two columns.

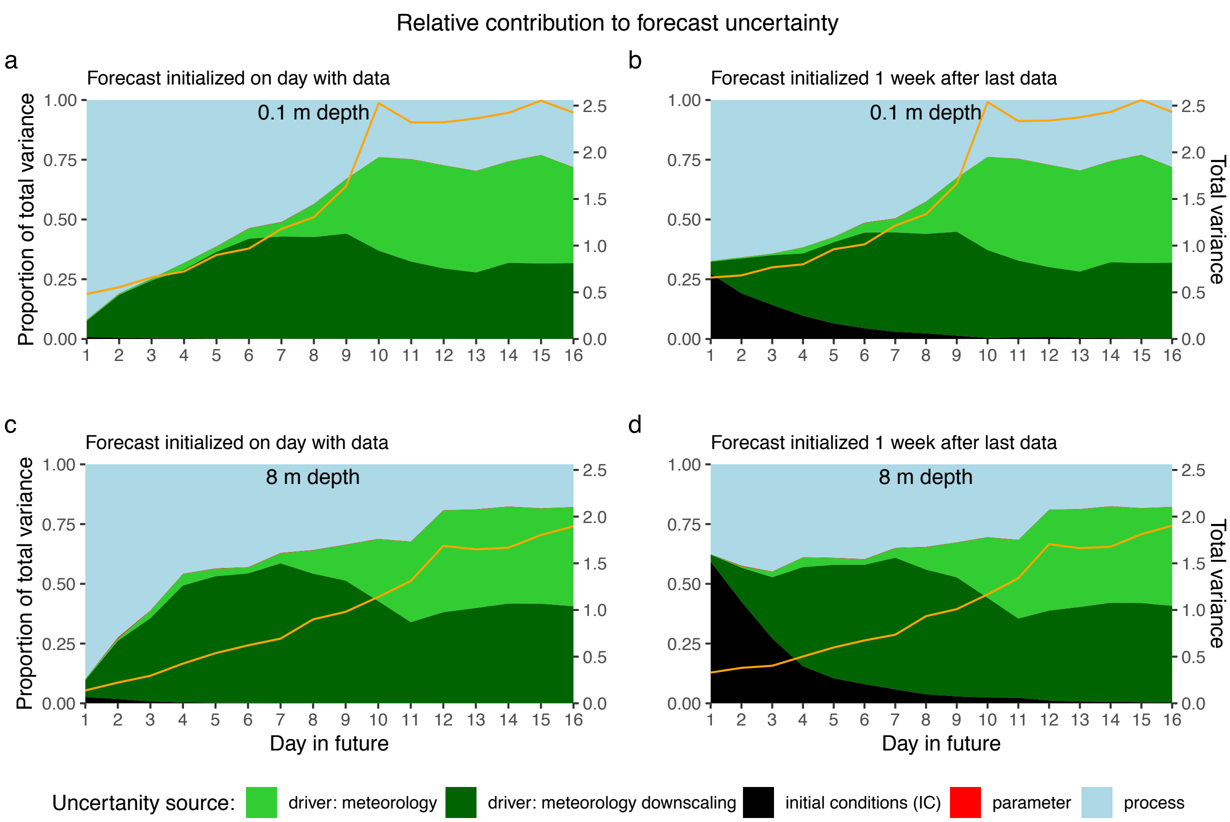

Supporting Information Figure 4. The relative contributions of the individual sources of uncertainty (left axis) to the total forecast uncertainty (right axis, orange line) varies through the 16-day forecast horizon. This forecast was initialized on 1 December 2018 and is one of three 16-day forecasts (with Supporting Information Figure 2 and Supporting Information Figure 3) that were averaged to create Figure 6. Two depths are shown (0.1 m – a, b; 8.0 m – c, d) and the relative contributions of initial condition uncertainty without (left) and with (right) gaps in water temperature sensor observations are shown in the two columns.