

The role of predictive uncertainty in the operational management of reservoirs

E. TODINI

BiGeA, University of Bologna, Via Zamboni, 67–40126 Bologna, Italy
ezio.todini@unibo.it

Abstract The present work deals with the operational management of multi-purpose reservoirs, whose optimisation-based rules are derived, in the planning phase, via deterministic (linear and nonlinear programming, dynamic programming, etc.) or via stochastic (generally stochastic dynamic programming) approaches. In operation, the resulting deterministic or stochastic optimised operating rules are then triggered based on inflow predictions. In order to fully benefit from predictions, one must avoid using them as direct inputs to the reservoirs, but rather assess the “predictive knowledge” in terms of a predictive probability density to be operationally used in the decision making process for the estimation of expected benefits and/or expected losses. Using a theoretical and extremely simplified case, it will be shown why directly using model forecasts instead of the full predictive density leads to less robust reservoir management decisions. Moreover, the effectiveness and the tangible benefits for using the entire predictive probability density instead of the model predicted values will be demonstrated on the basis of the Lake Como management system, operational since 1997, as well as on the basis of a case study on the lake of Aswan.

Key words predictive knowledge; predictive uncertainty; reservoir management

INTRODUCTION

Most of the large reservoirs in the world are currently managed either on the basis of politically influenced empirical rules or on the basis of optimisation-based rules derived either via deterministic (linear and nonlinear programming, dynamic programming, etc.) or stochastic (generally stochastic dynamic programming) approaches aimed at maximising the long-term reservoir management benefits.

The derived, deterministic or stochastic management rules, are then triggered in operation on the basis of future inflow predictions, mostly provided by seasonal or shorter-term hydro-meteorological forecasts. The type of the required prediction may vary according to the scope of the reservoir, with the reservoir size as well as with the inherent situation. For example, when a flood is bound to happen, the decision maker may become risk averse, thus changing the primary objective from benefit optimisation to flood risk management. This implies making use of shorter-term (weekly, daily or even hourly) operational flood forecasting systems.

Today, most of these forecasts, although inherently highly uncertain, are substantially treated as deterministic in the decision making process. This inevitably results in strongly sub-optimal management, which can in fact be definitely (and simply) improved by accounting for and incorporating the “predictive knowledge” in the decision making process.

Whenever a forecasting model is used, in order to take informed decisions, it is common practice to directly route the forecasted volumes into the reservoir for assessing via simulation the expected future status of the storage, together with the potential effects of decisions.

This approach, which in this paper will be referred to as “deterministic”, in that it implicitly assumes no uncertainty in the models’ forecasts, has only marginally improved reservoirs’ operations, apart from rare cases where extremely accurate and reliable predictive models are used for very short forecasting lead times.

Using a theoretical and extremely simplified case, it will be shown why neglecting the full extent of predictive knowledge leads to less reliable reservoir management. Moreover, the effectiveness and the tangible benefits for using the measure of predictive knowledge (the entire probability density) as opposed to the actual predicted value, will be demonstrated on the basis of the Lake Como management system, which has been operational since 1997, as well as on the basis of a case study on the lake of Aswan.

A SIMPLE CASE

In the last two decades, when dealing with forecasting, and in particular with real time flood forecasting, interest has grown on the concept and use of “predictive uncertainty”, but with limited benefits, given that it has always been perceived in the negative way of an expression of “lack of knowledge”.

As in the classical case of the half-empty/half-full glass, one must realise that what is generally referred to as a measure of “predictive uncertainty” must be more effectively interpreted as a measure of “predictive knowledge” (Fig. 1). And it is the measure of predictive knowledge, by incorporating it in the decision making process, that allows taking educated decisions.

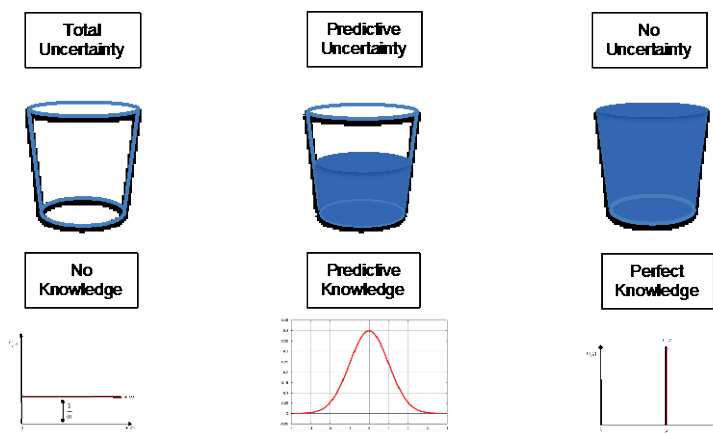


Fig. 1 The representation of predictive uncertainty and predictive knowledge.

When dealing with measurable quantities, such as for instance water volumes, water levels, etc., one of the most appropriate ways of describing our understanding (knowledge) on their potential occurrence is by using a probability density function. While the total ignorance of a quantity can be mathematically represented by a uniform density lying over an extremely wide range (for example, from 0 to $+\infty$) and the perfect knowledge can be described by a Dirac delta function (Dirac, 1958, pp. 58–61) over the exact value, any intermediate situation of knowledge can be mathematically represented by an appropriate probability density function, which can be both interpreted as a measure of uncertainty or as a measure of knowledge.

Having established a set of operating rules in the planning phase, it is reasonable to assume that if one had the perfect knowledge of a future occurrence then, given the established operating rules, he could choose the most appropriate management strategy. Unfortunately, apart from very special and practically meaningless cases, it will never be possible to reach the full knowledge of future occurrences. Consequently, it is fundamental to: (a) correctly assess a realistic measure of the available knowledge on the future events, and (b) make full use of this measure of knowledge to maximise the likelihood of taking the correct decisions.

The reasons for using the predictive knowledge in the decision making process can be illustrated using the following simple theoretical schematic example.

The reservoir in Fig. 2 is operated for irrigation and energy production purposes. Operating rules have been established on a long-term basis, from which the average value L_1 of water stored in the reservoir during the filling period is established to be:

$$L_1(X) = c_1 X \quad (1)$$

where X (m^3) is the volume of water and c_1 ($\text{€}/\text{m}^3$) the unit cost per unit volume of water stored.

A flood wave has been forecast which will raise the level of the reservoir from the current level to a future value, substantially higher than the reservoir top. Losses will occur if the water volume becomes larger than the maximum reservoir content and the economic losses can be expressed as L_2 , a damage function of the exceeding volume ($V - V_{\max}$) (assumed to be linear for simplifying the calculation as well as the reasoning without loss of generality).




Fig. 2 The simple schematic representation of the reservoir management problem.

$$\begin{cases} L_2(V) = 0 & V \leq V_{\max} \\ L_2(V) = c_2 (V - V_{\max}) & V > V_{\max} \end{cases} \quad (2)$$

where $(V - V_{\max})$ (m^3) is the volume of water in excess, V_{\max} the maximum content in the reservoir and c_2 ($\text{€}/\text{m}^3$) the unit cost of damages per unit of water volume in excess. It is also reasonable to assume that $c_2 \gg c_1$, to account for the consequences of flooding.

As previously stated, this problem is purposely simplified in order to enlighten the benefits of using the full measure of predictive knowledge instead of merely the model forecast.

Let us first assume that a future volume model forecast \hat{V} is available and it is our intention to use this quantity for establishing how much water should be preventively released in order to minimise the overall losses. In order to find the optimal release V_x we must then minimise the following function expressing the total losses:

$$\min_{V_x} L(V_x) = L_1(V_x) + L_2(\hat{V}) = c_1 V_x + \begin{cases} 0 & \forall \hat{V} - V_x \leq V_{\max} \\ c_2 (\hat{V} - V_{\max} - V_x) & \forall \hat{V} - V_x > V_{\max} \end{cases} \quad (3)$$

The minimum (being $c_1 \ll c_2$) is obviously reached when $V_x = \hat{V} - V_{\max}$, namely when the stored water volume equals the maximum reservoir content.

If \hat{V} is a “perfect” forecast, namely the future value V will be identical to the predicted value \hat{V} , (which implies assuming a Dirac delta over \hat{V} as the measure of our predictive knowledge) then this rule represents the best possible choice. But, as all would agree, \hat{V} cannot be perfect and it is unlikely that it will equal the actual occurrence of the future volume V .

Conversely, let us imagine that, from a hindcast assessment of performances of the available forecasting model one is able to correctly assess the “predictive density”, namely the probability density which expresses our best knowledge on the occurrence of the future value V once we are aware of what the forecasting model prediction \hat{V} is.

In order to simplify the problem without losing generality, let us assume that this predictive density, expressing our perception of the probability of occurrence of a future value once we are aware of the model forecast, can be represented by a Gaussian probability density with the following characteristics:

$$f(V|\hat{V}) = \frac{e^{-\frac{(V - \mu_{V|\hat{V}})^2}{2\sigma_{V|\hat{V}}^2}}}{\sqrt{2\pi\sigma_{V|\hat{V}}^2}} \quad (4)$$

where $\mu_{V|\hat{V}} = E\{V|\hat{V}\}$ is the expected conditional mean of the future value given the model forecast and $\sigma_{V|\hat{V}}^2 = \text{Var}\{V|\hat{V}\}$ the variance of the departures of the true unknown value from the expected conditional mean. As mentioned earlier, if correctly assessed, this predictive density is the best representation of our knowledge on the future occurrence. The conditional variance $\sigma_{V|\hat{V}}^2$ is indicative of the quality of our knowledge. In the case of an utopic “perfect” predictive model $\sigma_{V|\hat{V}}^2$ would tend to zero with the probability density degenerating into a Dirac delta over the value $V = \hat{V}$.

The effect of preventively releasing a volume of water V_x from the reservoir, will modify this conditional density as follows:

$$f(V|\hat{V}, V_x) = \frac{e^{-\frac{[V-(\mu_{V|\hat{V}}-V_x)]^2}{2\sigma_{V|\hat{V}}^2}}}{\sqrt{2\pi\sigma_{V|\hat{V}}^2}} \quad (5)$$

While the variance will not be affected by the released volume V_x , the mean of the predictive density, representing the future final volume stored in the reservoir, will become $\mu_{V|\hat{V}} - V_x$, namely the expected future volume $E\{V|\hat{V}\}$ conditional on the model forecast minus the released volume V_x . The variance of the predictive density is unchanged because no additional information (knowledge) on the future event, has been provided.

The optimal release can be found by minimizing the expected losses with respect to V_x , as:

$$\min_{V_x} E\{L(V_x)\} = c_1 V_x + c_2 \int_{V_{\max}}^{+\infty} (V - V_{\max}) \frac{e^{-\frac{[V-(\mu_{V|\hat{V}}-V_x)]^2}{2\sigma_{V|\hat{V}}^2}}}{\sqrt{2\pi\sigma_{V|\hat{V}}^2}} \quad (6)$$

where the expectation is only taken over the second term, since V_x , is not uncertain: it is a fixed quantity, once the decision is taken.

After a series of algebraic manipulations, the following interesting result can be obtained:

$$V_x = \mu_{V|\hat{V}} - V_{\max} + \sigma_{V|\hat{V}} N^{-1} \left\{ 1 - \frac{c_1}{c_2} \right\} \quad (7)$$

where $N^{-1} \left\{ 1 - \frac{c_1}{c_2} \right\}$ represents the value of a Standard Normal variable with probability $1 - \frac{c_1}{c_2}$.

The analysis of this result reveals that releasing $V_x = \mu_{V|\hat{V}} - V_{\max}$ is best only in two cases. Either when the unit losses for releasing water c_1 are identical to the unit cost of damages c_2 (which is not so, since the interesting case is when $c_2 \gg c_1$), or in the case of a perfect model forecast (namely, $V = \mu_{V|\hat{V}} = \hat{V}$ and $\sigma_{V|\hat{V}} = 0$) which again leads to $V_x = \hat{V} - V_{\max}$. In all other cases, one must release more than just $V_x = \mu_{V|\hat{V}} - V_{\max}$. The reason why is very simple. Not being entirely certain of what may occur, one must use the precaution principle. The larger the uncertainty is, the more cautious one must be and preventively release a larger amount of water than what might be assumed optimal under the hypothetical assumption of perfect knowledge of the future.

REAL WORLD EXAMPLES

The benefits arising from the use of the “predictive knowledge” concept in the operational management of a reservoir can be shown with two examples. The first relates to the natural lake of Como and the second to the High Aswan Dam reservoir. Note that, as opposed to the previous extremely simplified theoretical example based on two trade-off costs c_1 and c_2 , both real cases were optimized using more complex objective functions involving environmental, flooding and energy production costs.

Lake Como

Lake Como in northern Italy is mainly used for irrigation and hydro-electrical power production. Its outlet is controlled through a dam at Olginate and management of the lake has to cope with the necessity of saving water to satisfy agricultural and hydro-electrical demands and, at the same time, must guarantee safety against flooding, the risk of which has increased in the past 40 years.

The management problem is complicated by the relatively small reservoir control volume (246.5 Mm³, roughly 1/20 of the yearly inflow volume combined to the reduced hydraulic capacity of the downstream gates that allow releasing 900–1000 m³s⁻¹ whereas the inflow reaches 1800–2000 m³s⁻¹, which leads to rapid rising and filling of the lake (3–5 days) in the case of large flood events.

The two basic requirements of management were: optimisation of water resources availability as a long-term objective and flood risk management as a short-term objective. To cope with the long-term objectives, the operating rules were established on a 10-day basis using a Stochastic Dynamic Programming (SDP) algorithm (Todini 1999). In order to cope with floods, these 10-day rules were then taken as one of the targets for deriving a shorter-term operating rule, the second objective being flood crest subsidence. This shorter-term rule was derived conditionally to daily inflow forecasts provided by a simple Nearest Neighbour (NN) based forecasting model (Yakovitz 1987). Using the NN approach, a conditional predictive density was estimated and used in the decision making process for the estimation of the expected losses to be minimised (Todini 1999). The effect of the resulting operating rule was simulated over a period of 15 years (1981–1995) and compared to the recorded results deriving from the actual lake management. The obtained improvement was outstanding:

- (1) The frequency of the city of Como flooding was reduced by over 30%;
- (2) The water deficit was reduced on average by 110 Mm³/year (12%);
- (3) The electricity production was increased by 3%.

A decision support system (DSS) was operationally installed in 1997 and constantly used by the lake manager since. A recent evaluation of Lake Como management during a period of drought years 2000–2006 revealed that the management simulated by the DSS and the actually performed management based on the suggested releases are not only very close, but the real management produced better results than the simulated ones. This is due to the availability of additional information, such as meteorological forecasts, satellite imagery of cloud cover, etc., which can support the lake manager in taking the final decision.

The Reservoir of Aswan

In 1968, Egypt completed the High Aswan Dam (HAD) and thus created a powerful tool to master the Nile. The HAD has an active storage volume of 106 Billion m³. The importance of the HAD is based on its capability to reliably supply Egypt (and Sudan) with water for irrigation purposes throughout the year and, with decreasing importance, to produce hydroelectric energy. The medium annual flow measures 85 Billion m³, the actual annual irrigation demand of Egypt being of 54.2 Billion m³ and the Sudan abstraction of 16.5 Billion m³ (data from HAD control centre, Cairo).

CONCLUSIONS

There are three basic concluding remarks to be drawn from the present work. The first is the fact that in order to promote the use of “predictive uncertainty” among the decision makers in their decision-making process, it is essential to communicate it by looking at the full half of the glass. In other words what is usually referred to as “predictive uncertainty” must be more effectively delivered as “predictive knowledge” to mark the fact that more than a negative measure of what we do not know, it is a positive measure of what we do know. The second conclusion is that the use of the full predictive density in reservoir management always leads to expected results, which are more reliable and more beneficial than the ones obtained via empirical approaches or based on optimised rules directly triggered by the model forecasts. The third conclusion is that in order to make the stakeholders aware of the improvements they may expect from using the full predictive density it is essential to let them compare, over the historical records, the different performances obtained either by taking decisions according to their previous rules or by using, as discussed in this paper, the information on the predictive knowledge of the future events, embedded in the full predictive density.

REFERENCES

- Dirac, P. A. M. (1958) *The Principles of Quantum Mechanics* (4th ed.), International Series of Monographs on Physics 27, Oxford University Press, New York (USA), pp. 328. ISBN 978-0-19-852011-5.
- Todini E. (1999) Using phase-space modelling for inferring forecasting uncertainty in non-linear stochastic decision schemes, *Journal of Hydroinformatics*, 01.2, 75 – 82.
- Yakovitz, S. (1987) Nearest neighbour methods for time series analysis. *J. Time Series Analysis* 8(2), 235–247.