Environmental Industry Letters
https://journals.tultech.eu/index.php/eil
ISSN: 2806-2965
Volume 4, Issue 1
DOI: https://doi.org/10.15157/eil.2026.4.1.1-20
Received: 30.01.2026; Revised: 05.03.2026; Accepted: 10.03.2026

Research Article

A Genetic-Algorithm-Based Multi-Objective Decision Support System for Cascade Reservoir Water Resource Management

Aleksander Vokhmintcev[1*], Tatyana Vokhmintceva[2], Natalia D. Ivanova[2], Hussein Alkattan[3], Mostafa Abotaleb[2], Mustafa Adel[4], Ioannis Adamopoulos[5]

[1] Chelyabinsk State University, Chelyabinsk, Russia

[2] Engineering School of Digital Technologies, Yugra State University, Khanty-Mansiysk, Russia

[3] Department of System Programming, South Ural State University, Chelyabinsk, Russia

[4] Modern Academy for engineering and Technology, Cairo, Egypt

[5] Department of Public Health and Policies, Hellenic Open University, Greece

*** Abotalebmostafa@bk.ru**

Abstract

Growing volumes of publicly available hydrological, meteorological, and reservoir-operation data create new opportunities for evidence-based water resource management. However, converting these data into actionable operating strategies requires optimisation frameworks capable of addressing the non-linear, non-convex, and conflicting objectives of reservoir operation. This study develops and evaluates a genetic-algorithm-based decision support system (DSS) for a three-reservoir cascade serving irrigation, municipal water supply, hydropower generation, and flood control. A real-coded, constraint-handling Non-dominated Sorting Genetic Algorithm II (NSGA-II) is implemented in Python and coupled with a mass-balance simulation model driven by a ten-year stochastically generated monthly inflow record calibrated to a monsoon-influenced river basin. The optimisation considers a 45-parameter cyclic hedging-rule policy, producing implementable and repeatable management strategies rather than fixed release schedules. Eight independent runs (150 individuals, 300 generations) produce a consolidated Pareto front of 417 non-dominated policies balancing water-supply deficit, hydropower generation, and flood-control performance. A TOPSIS-based compromise-selection layer converts these solutions into operating recommendations under three managerial preference profiles. Compared with a conventional Standard Operating Policy baseline, the balanced GA-optimised policy increases hydropower generation from 6,453.9 to 6,564.2 GWh over ten years and reduces the flood-control exceedance index by 90%, while the water-supply-priority policy increases generation to 6,789.3 GWh. The results demonstrate a deliberate reliability-resilience-vulnerability trade-off, with GA policies rationing water more frequently but moderately, thereby redistributing risk across the cascade. Sensitivity and convergence analyses further confirm the robustness of the optimisation framework.

Keywords: Genetic algorithm; NSGA-II; Reservoir operation; Water resources management; Decision support system; Multi-objective optimisation

INTRODUCTION

Reservoir systems lie at the centre of modern water resource management, simultaneously serving irrigation and municipal water supply, hydropower generation, flood control and, increasingly, environmental flow requirements. The task of deciding how much water to release from each reservoir in a system, in every period, so that these often-conflicting purposes are served as well as possible, is a classical yet still unresolved problem in water resources systems analysis. The difficulty is compounded when several reservoirs are hydraulically connected in a cascade, because a release decision at an upstream reservoir propagates downstream and alters the feasible operating space of every reservoir below it.

The growing availability of large, systematically archived hydrological and water-resources datasets -- long-record streamflow archives such as the Global Runoff Data Centre (GRDC) database and national systems such as the United States Geological Survey's National Water Information System (NWIS) [17,18], together with satellite-derived storage-change, precipitation and land-surface products -- has made increasingly detailed, data-driven characterisation of basin hydrology routine. Yet a large hydrological dataset only becomes managerially useful once it is embedded in a decision-making framework that can translate historical or synthetic inflow information into concrete, implementable operating advice. Reservoir operation optimisation is exactly that translation layer, and this study is concerned with building, rigorously testing, and packaging such a framework as a managerial decision support system (DSS).

Reservoir operation problems are notoriously difficult for classical gradient-based or linear/dynamic-programming optimisation methods because the governing physical relationships -- storage-area-elevation curves, turbine power equations, and hedging-type operating rules with discontinuous switching conditions -- are highly non-linear, and because realistic formulations involve multiple, non-commensurable, conflicting objectives for which no single scalar optimum exists. Evolutionary algorithms, and genetic algorithms (GAs) in particular, are attractive for this class of problem because they require no gradient information, tolerate non-smooth and discontinuous objective functions and constraints, and naturally generate a population of candidate solutions that can be post-processed into a full Pareto-optimal trade-off surface rather than a single point solution [1,2]. Since the foundational applications of GAs to water resources problems, a substantial literature has established their effectiveness for reservoir operation, and multi-objective variants such as the Non-dominated Sorting Genetic Algorithm II (NSGA-II) have become a standard tool for deriving Pareto-optimal reservoir operating policies [3,4].

This study makes four contributions. First, it implements every component of a constrained, multi-objective NSGA-II optimiser (non-dominated sorting, crowding distance, simulated binary crossover, polynomial mutation and constrained-domination tournament selection) directly in Python, without reliance on an external evolutionary-computation library, so that the complete algorithmic mechanics of the decision support system are transparent and auditable by a water utility's own engineering staff. Second, it couples this optimiser to a physically explicit, vectorised mass-balance simulation model of a three-reservoir cascade system -- including evaporation, storage-dependent hydropower head, and a genuine storage-feedback hedging-rule operating policy -- driven by a stochastically generated, ten-year monthly inflow record calibrated to the statistical character of a real monsoon-influenced river basin. Third, rather than optimising an open-loop, non-repeatable release schedule, the decision variables are structured as a compact cyclic operating policy (target releases plus hedging and flood-release parameters for each reservoir and calendar month), so that the optimisation output is directly implementable as a standing operating rule, comparable on equal terms with conventional practice. Fourth, the resulting Pareto front is embedded in a TOPSIS-based compromise-selection layer that allows a manager to translate stated preferences among water supply, hydropower and flood-control objectives directly into a recommended operating policy, completing the loop from optimisation algorithm to managerial decision support tool.

The remainder of this paper is organised as follows. Section 2 reviews prior applications of genetic and other evolutionary algorithms to reservoir operation. Section 3 describes the case-study system, the simulation model, the multi-objective optimisation formulation, and the NSGA-II decision support system. Section 4 reports the optimisation results, including convergence behaviour, parameter sensitivity, Pareto-optimal trade-offs and a detailed comparison against the baseline policy. Section 5 presents the decision support system itself and discusses its managerial use and limitations. Section 6 concludes.

LITERATURE REVIEW

Genetic algorithms, introduced as a general adaptive search procedure modelled on natural selection [1] and subsequently formalised and popularised for search and optimisation problems [2], were among the first metaheuristics applied to reservoir operation because of their ability to handle non-convex, non-differentiable objective functions and mixed discrete-continuous decision spaces. Early applications demonstrated that real-coded GAs could reproduce, and in many cases improve upon, operating rules obtained from stochastic dynamic programming for single- and multi-reservoir systems, while avoiding the curse of dimensionality that limits dynamic programming to systems with only a few state variables [5,6]. Wardlaw and Sharif [6] provided one of the most influential early evaluations of GA design choices -- coding scheme, crossover and mutation operators, and elitism -- for a four-reservoir operation problem, concluding that real-value coding with appropriately tuned operators outperformed binary coding on this class of problem, a finding consistent with the real-coded representation adopted in the present study.

A parallel line of work extended single-objective GA reservoir operation to explicitly long-term, multi-year operating horizons. Chen [7] applied a real-coded GA to derive long-term operating rules for a single reservoir and reported that real coding converged more reliably than binary coding for continuous release decision variables, reinforcing the choice of a real-valued chromosome representation here. As reservoir operation problems were increasingly recognised as inherently multi-objective -- water supply, hydropower, flood control and environmental flows rarely admit a common scalar objective without ad hoc weighting -- attention shifted toward Pareto-based multi-objective evolutionary algorithms. Reddy and Kumar [8] used a multi-objective particle swarm optimiser to generate Pareto-optimal trade-offs for reservoir operation and explicitly framed the resulting trade-off curve as a decision support tool for reservoir managers, a framing this study adopts and extends with a GA-specific, TOPSIS-based compromise-selection mechanism.

Two comprehensive review articles frame the current state of the art. Labadie [5] reviewed optimal operation of multireservoir systems and concluded that evolutionary and other heuristic search methods had become competitive with, and in many practical settings preferable to, classical stochastic dynamic programming and linear/non-linear programming approaches, particularly once realistic non-linearities and multiple objectives are introduced. Nicklow et al. [3], reporting on behalf of the ASCE Task Committee on Evolutionary Computation in Water Resources Engineering, documented the breadth of GA applications across water distribution design, groundwater management and reservoir operation, and identified hybridisation with local search, improved constraint-handling and integration into operational decision support as priority research directions -- directions this study addresses directly through its constrained-domination tournament selection [11] and its TOPSIS-based decision layer [13]. More recently, Maier et al. [4] surveyed the broader use of evolutionary algorithms and other metaheuristics across the water resources field and emphasised that translating multi-objective optimisation results into operational decisions -- rather than treating the Pareto front as an academic end-point -- remains one of the field's central practical challenges, which is the specific gap this study's managerial decision support layer is designed to close.

The multi-objective genetic algorithm used in this study, NSGA-II [9], together with its simulated binary crossover operator [10] and Deb's constrained-domination tournament for handling inequality constraints without explicit penalty-weight tuning [11], has become one of the most widely adopted multi-objective evolutionary algorithms in water resources engineering because of its computational efficiency, elitism, and explicit diversity preservation via crowding distance. Its convergence and diversity properties are commonly assessed using indicators such as the hypervolume, originally proposed alongside the Strength Pareto Evolutionary Algorithm [16], which this study also adopts to track and compare convergence across independent GA replicates and across the parameter sensitivity study. Complementing the optimisation itself, Hashimoto, Stedinger and Loucks [12] established the reliability-resilience-vulnerability (RRV) framework that remains the standard basis for evaluating reservoir system performance beyond a single aggregate objective, and is used in Section 4 of this study to unpack the trade-offs embedded in the optimised Pareto front. Finally, converting a Pareto front into a single operational recommendation is itself a multi-criteria decision problem; this study uses the TOPSIS method [13], a widely used compromise-ranking technique, to let a manager's stated preference weights select a specific operating policy from the optimised trade-off surface.

MATERIALS AND METHODS

Case-Study System and Data

The case-study system is a cascade of three multi-purpose reservoirs (Reservoir 1, upstream; Reservoir 2, mid-cascade; Reservoir 3, downstream), each supplying a combined irrigation-and-municipal off-take and a hydropower plant, and each subject to a flood-control storage limit set at 90% of full supply capacity (Table 1). Reservoir outflow not diverted for consumptive use continues downstream as inflow to the next reservoir in the cascade, together with the incremental local (ungauged) inflow generated within that reach.

Local monthly inflow to each of the three sub-catchments was generated using the classical lag-one, seasonal log-normal streamflow generation model of Thomas and Fiering [14], as implemented in standard water resources systems planning practice [15]. Seasonal mean inflow, coefficient of variation, and a lag-one serial correlation coefficient of 0.30 were parameterised to reproduce the strong wet-season/dry-season contrast (peak inflow in June-September, low flow in December-March) and the coefficient-of-variation range (approximately 0.28-0.48, higher in the low-flow season) reported in the hydrological literature for monsoon-influenced river basins of the kind monitored in long-record archives such as GRDC and USGS NWIS [17,18]. A ten-year (120-month) synthetic record was generated independently for each of the three sub-catchments with a fixed random seed, giving a reproducible, internally consistent hydrological forcing sequence against which every operating policy in this study -- the Standard Operating Policy baseline and every GA-generated candidate -- is evaluated identically. This case-study design follows common practice in the reservoir operation optimisation literature of using a statistically representative synthetic hydrological sequence (rather than a single historical record, which by construction provides only one realisation of basin hydrology) to test candidate operating policies under a broader range of wet, average and dry conditions than any single historical record could provide.

Table 1. Physical and operational parameters of the three-reservoir cascade case-study system.

Parameter	Reservoir 1 (upstream)	Reservoir 2 (mid-cascade)	Reservoir 3 (downstream)
Full supply capacity, S_{max} (MCM)	900.0	650.0	1200.0
Dead storage, S_{dead} (MCM)	90.0	65.0	120.0
Initial storage, S_0 (MCM)	500.0	350.0	700.0
Flood-control threshold, $0.90*S_{max}$ (MCM)	810.0	585.0	1080.0
Minimum release, R_{min} (MCM/month)	5.0	5.0	5.0
Maximum release, R_{max} (MCM/month)	420.0	300.0	520.0
Consumptive diversion fraction, f_{div}	0.4	0.3	0.5
Turbine efficiency, eta	0.88	0.85	0.9
Minimum net head, H_{min} (m)	25.0	20.0	30.0

Maximum net head, H_{max} (m)	65.0	48.0	75.0
Installed capacity (MW)	120.0	80.0	160.0

Reservoir System Simulation Model

Each candidate operating policy is evaluated by a fully vectorised (population-wide) monthly mass-balance simulation of the cascade. Reservoir storage evolves according to

$$S_{i,t+1} = S_{i,t} + I_{i,t} - E_{i,t} - R_{i,t} - Sp_{i,t} \tag{1}$$

where $I_{i,t}$ is the sum of local inflow and the routed, non-diverted release from the upstream reservoir (for Reservoir 1, local inflow only), $Sp_{i,t}$ is spill occurring whenever the post-release storage would otherwise exceed full supply capacity, and evaporation is computed from a storage-dependent reservoir surface area following a standard power-law elevation-storage-area relationship,

$$E_{i,t} = e_{m(t)} \cdot a_i S_{i,t}^{b_i} \tag{2}$$

with $e_{m(t)}$ the mean monthly net evaporation depth for calendar month m(t), and a_i, b_i the reservoir-specific area-storage coefficients (Table 1). The actual release delivered in each period is limited to the water physically available above dead storage after evaporation, so that an infeasible decision release is never applied outright but is instead automatically rationed by the simulation itself, consistent with real reservoir operation.

Operating-Policy Decision Variables

A first version of this study encoded the decision space as one free release value per reservoir per month of the full ten-year simulation horizon (360 real-valued genes). This open-loop encoding proved poorly matched to evolutionary search: with no structural link between genes, undirected recombination struggled to discover the specific, highly structured 360-dimensional target implied by the demand pattern, and -- more importantly -- the resulting "policy" would not have been implementable by a real operator beyond the single synthetic hydrological sequence used to derive it. Following standard reservoir-operation practice, in which operating rules are expressed as reusable, storage-feedback rule curves or hedging rules rather than one-off release schedules, this study instead searches over a compact, cyclic hedging-rule operating policy for each reservoir, applied identically in every year of the simulation and hence genuinely repeatable as an operating rule.

For each reservoir i and calendar month m, the decision (target) release under normal storage conditions is $T_{i,m}$. If storage falls below a hedging-trigger level S_i^{hedge}, defined as a fraction h_i of usable (above-dead) storage,

$$S_i^{hedge} = S_i^{dead} + h_i(S_i^{max} - S_i^{dead}) \tag{3}$$

the release is rationed according to

$$R_{i,t} = T_{i,m(t)} \left(\frac{S_{i,t} - S_i^{dead}}{S_i^{hedge} - S_i^{dead}} \right)^{k_i} \tag{4}$$

where k_i is a GA-optimised hedging-severity exponent controlling how aggressively release is cut as storage approaches dead storage. Independently, whenever storage exceeds the flood-control threshold S_i^{flood}, additional flood-mitigation release $R_{i,t}^{flood}$, proportional to a GA-optimised flood-release gain g_i, is added to the release computed above,

$$R_{i,t}^{flood} = g_i \max(0,\ S_{i,t} - S_i^{flood}) \tag{5}$$

and the resulting release is finally clipped to the physical bounds [R_i^{min}, R_i^{max}] of the reservoir's outlet works. Each reservoir therefore contributes 15 decision variables (12 monthly targets $T_{i,1..12}$ plus h_i, k_i and g_i), for a total of 45 decision variables across the three-reservoir cascade -- an order of magnitude smaller than the open-loop encoding while representing a genuinely implementable operating rule. The Standard Operating Policy baseline against which the GA is compared (Section 3.5) is expressed as one particular, non-optimised parameter set within this same 45-parameter structure, which keeps the comparison between "current practice" and "GA-optimised practice" strictly like-for-like.

Hydropower Generation

Hydroelectric generation at each reservoir is computed from the standard power equation

$$P_{i,t} = \frac{\eta_i \rho g Q_{i,t} H_{i,t}}{10^6} \quad [\mathrm{MW}] \tag{6}$$

where $Q_{i,t}$ is the mean monthly discharge through the turbine (converted from the monthly release volume), $H_{i,t}$ is the net hydraulic head, linearly interpolated between the minimum head at dead storage and the maximum head at full supply level using the average of start- and end-of-period storage, η_i is plant efficiency, ρ is water density and g is gravitational acceleration. Generated power is capped at each plant's installed capacity and converted to monthly energy in GWh.

Multi-Objective Optimisation Formulation

Three objectives, all cast as minimisation targets, are optimised simultaneously. The first objective combines the severity and the frequency of water-supply shortfalls. Writing $d_{i,t}$ for the consumptive-demand deficit at node i in period t and $\bar{D}_i$ for the mean demand at that node,

$$f_1 = \sum_{i,t} \left(\frac{d_{i,t}}{\bar{D}_i} \right)^2 + \alpha \sum_{i,t} \delta(d_{i,t} > 0) \tag{7}$$

where $\delta(\cdot)$ is an indicator function equal to 1 when its argument is true and 0 otherwise, and α (set to 1.0 in this study) weights the relative importance of shortfall frequency against shortfall severity. This two-part specification is deliberate: minimising the sum of squared deficits alone rewards concentrating shortfalls into a few severe events, which can perversely reduce the classical, frequency-based reliability metric of Hashimoto et al. [12] even as the squared-deficit objective improves. Because a real-coded genetic algorithm requires no gradient information, it optimises a non-smooth, indicator-based term of this kind just as readily as a smooth one -- a practical advantage over gradient-based alternatives that is one of the reasons GAs remain attractive for reservoir operation [3]. The second objective is total hydropower energy generation, expressed in minimisation form as its negative,

$$f_2 = -\sum_{i,t} EN_{i,t} \tag{8}$$

where $EN_{i,t}$ is monthly generated energy (GWh) from Equation 6. The third objective penalises storage held above the flood-control threshold, squared and normalised by the available flood-storage buffer of each reservoir,

$$f_3 = \sum_{i,t} \left(\frac{\max(0,\ S_{i,t} - S_i^{flood})}{S_i^{max} - S_i^{flood}} \right)^2 \tag{9}$$

Finally, an end-of-horizon carry-over storage constraint prevents the optimiser from artificially inflating short-term releases (and hence f2) by draining the system toward the end of the ten-year simulation horizon. Writing S_i^0 for initial storage and $S_{i,N}$ for storage at the final period N, the aggregate constraint violation is

$$g_1 = \sum_i \frac{\max(0,\ 0.6\, S_i^0 - S_{i,N})}{S_i^0} \tag{10}$$

which requires each reservoir to retain at least 60% of its initial storage by the end of the simulation horizon; g1 = 0 denotes a feasible solution. The complete optimisation problem is therefore to find the 45-parameter policy vector $x = (T_{i,m}, h_i, k_i, g_i)$ that minimises [f1(x), f2(x), f3(x)] subject to g1(x) = 0 and the parameter bounds in Table 2.

Genetic-Algorithm-Based Decision Support System

The optimisation engine is a real-coded, constraint-handling implementation of NSGA-II [9], written from first principles in Python and NumPy without reliance on an external evolutionary-computation library, so that every algorithmic step is directly auditable. The implementation comprises five components. First, fast non-dominated sorting partitions a population into Pareto fronts using Deb's constrained-domination relation [11]: a feasible solution always dominates an infeasible one; between two infeasible solutions, the one with smaller aggregate constraint violation dominates; and between two feasible solutions, standard Pareto dominance over (f1, f2, f3) applies. This removes the need for a hand-tuned penalty weight on the constraint term. Second, crowding-distance assignment estimates the local solution density around each individual within its front, used to preserve diversity along the trade-off surface. Third, binary tournament selection chooses parents by comparing (rank, crowding distance) pairs. Fourth, offspring are generated by simulated binary crossover (SBX) [10] with distribution index $\eta_c = 15$ and crossover probability $p_c = 0.9$, followed by polynomial mutation with distribution index $\eta_m = 20$ and per-gene mutation probability $p_m = 1/n_{var}$. Fifth, an elitist environmental selection step merges the parent and offspring populations, re-applies non-dominated sorting, and fills the next generation front-by-front, using crowding distance to break ties within the last admitted front. Both the $O(N^2)$ non-dominated-sorting and crowding-distance steps are implemented as fully vectorised NumPy array operations rather than nested Python loops, which was essential for making the roughly 1.4 million candidate-policy evaluations performed across this study's main experiment and sensitivity analysis computationally tractable.

Table 2. NSGA-II algorithm configuration used in the main experiment.

Parameter	Value
Decision variables	45
Population size	150
Number of generations	300
Crossover type	Simulated Binary Crossover (SBX)
Crossover probability, p_c	0.9
SBX distribution index, η_c	15.0
Mutation type	Polynomial mutation
Mutation probability, p_m	0.02222 ($1/n_{var}$)
Polynomial mutation index, η_m	20.0
Independent replicates (seeds)	8
Selection	Binary tournament (rank, crowding distance)
Constraint handling	Constrained-domination (Deb, 2000)

Eight independent optimisation replicates (distinct random seeds) were run, each with a population of 150 individuals evolved for 300 generations, to guard against reporting results specific to a single, possibly favourable, random initialisation. The final populations of all eight replicates (1,200 candidate policies) were pooled and re-sorted by non-dominated rank to obtain a single consolidated Pareto-optimal front, reported in Section 4, that represents the best trade-off surface identified across the entire experiment.

Baseline Standard Operating Policy

The GA-optimised policies are benchmarked against a conventional, practice-based Standard Operating Policy (SOP), of the kind widely used as the reference case in the reservoir-operation literature [5,6]. The SOP sets each month's target release to exactly meet consumptive demand ($T_{i,m} = D_{i,m} / f_{div,i}$), applies a hedging trigger at the bottom 15% of usable storage with linear rationing ($h_i = 0.15$, $k_i = 1.0$), and releases half of any storage held above the flood-control threshold as flood mitigation ($g_i = 0.50$) -- that is, the SOP is one specific, non-optimised parameter set within exactly the same 45-parameter hedging-rule structure that the GA searches over, ensuring that any performance difference reported in Section 4 reflects the optimisation itself rather than a difference in modelling framework.

Performance Metrics

Beyond the three optimisation objectives, candidate policies are also evaluated using the reliability, resilience and vulnerability (RRV) indices of Hashimoto et al. [12], computed independently at each reservoir node. Reliability is the fraction of periods in which demand is fully met; resilience is the conditional probability that a period of shortfall is followed by a period of full supply (a measure of recovery speed); and vulnerability is the average shortfall magnitude during failure periods, normalised by mean demand. Convergence of the NSGA-II optimisation is tracked using a Monte Carlo estimator of the hypervolume indicator [16] -- the volume of objective space dominated by the current population relative to a fixed reference point set 5% beyond the worst objective value observed anywhere in the study -- evaluated at every generation for every replicate so that both the mean convergence trajectory and its run-to-run variability can be reported.

Decision Support Compromise Selection

The consolidated Pareto front is converted into concrete operating recommendations using the Technique for Order of Preference by Similarity to Ideal Solution (TOPSIS) [13]. Given a manager-specified weight vector w = (w1, w2, w3) over the three objectives, each Pareto-optimal solution's objective vector is vector-normalised and weighted, and its closeness coefficient is computed as

$$C_p = \frac{d_p^-}{d_p^+ + d_p^-} \tag{11}$$

where d_p^+ and d_p^- are the Euclidean distances of solution p from the (weighted, normalised) ideal and anti-ideal points, respectively; the solution with the largest C_p is recommended under that preference profile. This study evaluates three illustrative managerial preference profiles -- a balanced profile (equal weight on all three objectives), a water-supply-priority profile (weight 0.6 on the deficit objective), and a hydropower-priority profile (weight 0.6 on the energy objective) -- to demonstrate how the decision support system responds to different stated priorities; in operational use, a manager would supply their own weights directly.

Computational Implementation

The complete framework -- stochastic hydrology generator, vectorised reservoir simulation, NSGA-II optimiser, baseline policy, TOPSIS compromise selection, hypervolume and RRV metrics, and all figures and tables in this manuscript -- is implemented in Python 3 using NumPy, SciPy, pandas and Matplotlib. All stochastic elements (hydrology generation, GA population initialisation, genetic operators, and Monte Carlo hypervolume estimation) use explicitly seeded NumPy random generators, making every reported result exactly reproducible.

RESULTS AND DISCUSSION

GA Convergence and Parameter Sensitivity

Across the eight independent NSGA-II replicates, the Monte Carlo hypervolume indicator increased monotonically on average and stabilised well before the 300th generation (Figure 1), indicating that the population had converged onto a stable trade-off surface rather than still actively improving when the run was terminated. The final-generation hypervolume averaged 2.370e+08 with a standard deviation of 6.876e+06 across the eight replicates -- a coefficient of variation of only 2.9%, indicating that the optimisation procedure is numerically stable and not strongly dependent on the random seed.

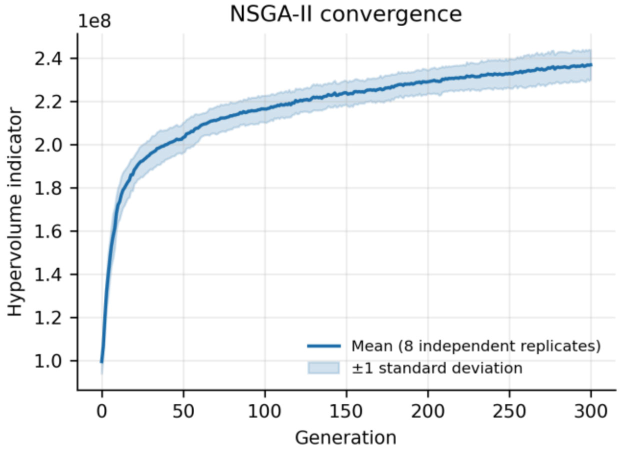

Figure 1. NSGA-II convergence: mean hypervolume indicator ± 1 standard deviation across 8 independent replicates (population 150, 300 generations).

A parameter sensitivity analysis (population size ∈ {60, 120, 200}; mutation-rate multiplier ∈ {0.5, 1.0, 2.0} × $1/n_{var}$; 3 seeds per combination; 150 generations) confirms the expected pattern that larger populations achieve higher final hypervolume for a fixed generation budget (Figure 2 and Table 3): mean hypervolume rose from 2.037e+08 at population 60 to 2.308e+08 at population 200, averaged across mutation settings. The mutation-rate multiplier had a comparatively minor effect within the tested range, with a mild optimum near the standard $1/n_{var}$ rate for smaller populations and near twice that rate for the largest population tested, consistent with the general principle that larger populations can profitably sustain slightly higher exploration pressure. On this basis, the main experiment's configuration (population 150, mutation probability $1/n_{var}$) sits close to the best-performing region of the tested parameter space while remaining computationally economical.

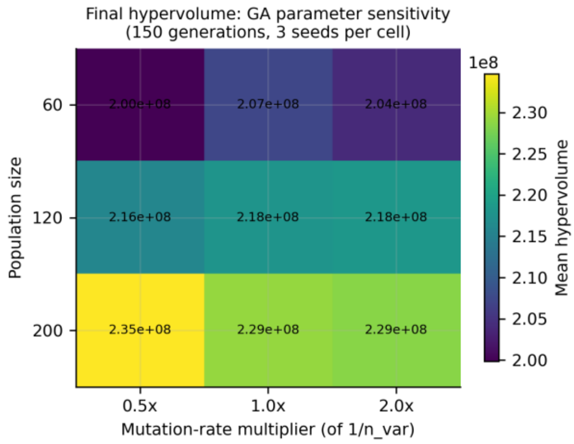

Figure 2. GA parameter sensitivity: mean final hypervolume as a function of population size and mutation-rate multiplier (150 generations, 3 seeds per cell).

Table 3. GA parameter sensitivity: final-generation hypervolume (mean and standard deviation across 3 seeds per cell).

Population size	Mutation-rate multiplier (x 1/n_{var})	Mean final hypervolume	Std. dev. final hypervolume
60	0.5	1.998e+08	4.982e+06
60	1.0	2.072e+08	6.983e+06
60	2.0	2.040e+08	6.328e+06
120	0.5	2.158e+08	9.032e+06
120	1.0	2.179e+08	7.870e+06
120	2.0	2.183e+08	9.768e+06
200	0.5	2.347e+08	3.632e+06
200	1.0	2.292e+08	9.545e+06
200	2.0	2.286e+08	1.790e+06

Pareto-Optimal Trade-Offs

Pooling the final populations of all eight replicates and re-applying non-dominated sorting yields a consolidated Pareto-optimal front of 417 distinct operating policies out of the 1,200 candidate policies evaluated in the final generation. Across this front, the water-supply deficit index f1 ranges from 101.3 to 309.9, ten-year hydropower generation ranges from 6,006.8 to 7,683.1 GWh, and the flood-control exceedance index f3 ranges from 1.0 to 78.0 -- a wide, well-populated trade-off surface (Figure 3) spanning policies that emphasise water-supply reliability, policies that emphasise energy generation, and policies that emphasise flood-control safety. As is typical of multi-objective reservoir optimisation, the front is markedly non-linear: the f1-f3 and f2-f3 projections in Figure 3 show that flood-control performance can be improved sharply for a modest sacrifice in water-supply deficit or energy generation over most of the front, but further gains in flood safety beyond roughly f3 ≈ 5 come at a steeply rising cost in the other two objectives -- information a manager can read directly from the trade-off curve when selecting an operating policy.

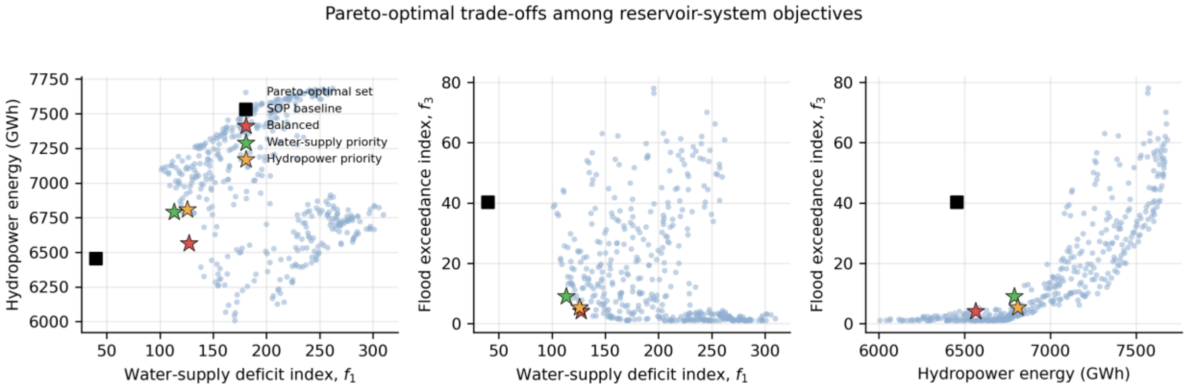

Figure 3. Pareto-optimal trade-offs among the water-supply deficit index (f1), hydropower energy generation, and the flood-control exceedance index (f3). The Standard Operating Policy baseline and the three TOPSIS-selected managerial recommendations (Section 5) are marked for reference.

Comparison with the Standard Operating Policy

Table 4 compares the Standard Operating Policy baseline against the three TOPSIS-selected GA-DSS recommendations introduced in Section 5. The SOP baseline achieves the lowest water-supply deficit of any policy examined in this study (967.3 MCM over the ten-year horizon, 93.9% period-wise reliability) because it is, by construction, a greedy demand-matching rule: it releases whatever is needed to meet consumptive demand whenever water is physically available, deferring to storage-depletion risk only through a narrow, fixed 15% hedging buffer. This greedy behaviour, however, provides no active flood-control response beyond releasing half of any storage surplus above the flood threshold, and the SOP's flood-exceedance index (f3 = 40.3) is far above any of the three GA-DSS recommendations (f3 between 4.1 and 9.0). The GA-optimised balanced policy simultaneously increases ten-year hydropower generation from 6,453.9 to 6,564.2 GWh (+1.7%) and reduces total spill from 1,430.3 to 266.1 MCM (-81%), while the water-supply-priority recommendation raises energy generation further, to 6,789.3 GWh.

Table 4. Ten-year performance comparison: Standard Operating Policy baseline versus GA-DSS recommendations under three managerial preference profiles.

Operating policy	**Total hydropower energy (GWh, 10 yr)**	**Total water-supply deficit (MCM, 10 yr)**	**Period-wise supply reliability (%)**	**Flood-control exceedance index f3 (-)**	**Total spill (MCM, 10 yr)**
Standard Operating Policy (SOP, baseline)	6,453.9	967.3	93.9	40.3	1,430.3
GA-DSS: Balanced preference	6,564.2	3,469.7	75.6	4.1	266.1
GA-DSS: Water-supply priority	6,789.3	3,089.8	77.8	9.0	320.8
GA-DSS: Hydropower priority	6,810.0	3,570.9	76.9	5.4	286.1

These gains are not obtained for free. Period-wise reliability under the GA-DSS policies (75.6-77.8%) is markedly lower than under the SOP baseline (93.9%). This is a genuine and, on reflection, theoretically expected outcome rather than a shortcoming of the optimisation: a hedging-based operating policy, by design, rations supply proactively and relatively often, in small increments, precisely so that it can avoid depleting storage during a prolonged dry sequence and can retain the flexibility to draw down surplus storage ahead of flood-prone months. A greedy, non-hedging policy such as the SOP instead supplies demand fully as long as water is physically available and only fails abruptly, and comparatively rarely, once storage is essentially exhausted. Section 4.5 examines this trade-off directly through the reliability-resilience-vulnerability framework and shows that it is consistent with established hedging theory rather than an artefact of the optimisation.

Reservoir Storage Dynamics

Figure 4 compares simulated storage trajectories for the GA-DSS balanced policy against the SOP baseline over the full ten-year case-study horizon. At Reservoir 1 (upstream), the GA-DSS policy visibly shaves the highest storage peaks -- most clearly around simulation years 2.7, 5.7 and 8.7, where the SOP trajectory presses against the full-supply capacity while the GA-DSS trajectory is held several tens of MCM below the flood-control threshold. At Reservoir 2 (mid-cascade), the effect is more pronounced still: the GA-DSS policy operates at a persistently lower average storage than the SOP baseline throughout the simulation, actively trading storage (and the water-supply reliability that storage buffers) for a wide flood-control margin that the SOP policy essentially never maintains. Reservoir 3 (downstream), which benefits from the largest catchment area and storage capacity in the cascade, shows the two policies tracking one another closely, with the GA-DSS trajectory only marginally lower on average -- indicating that most of the GA-optimised system's flood-control improvement is achieved by adjusting operation at the upstream and mid-cascade reservoirs rather than by suppressing storage system-wide.

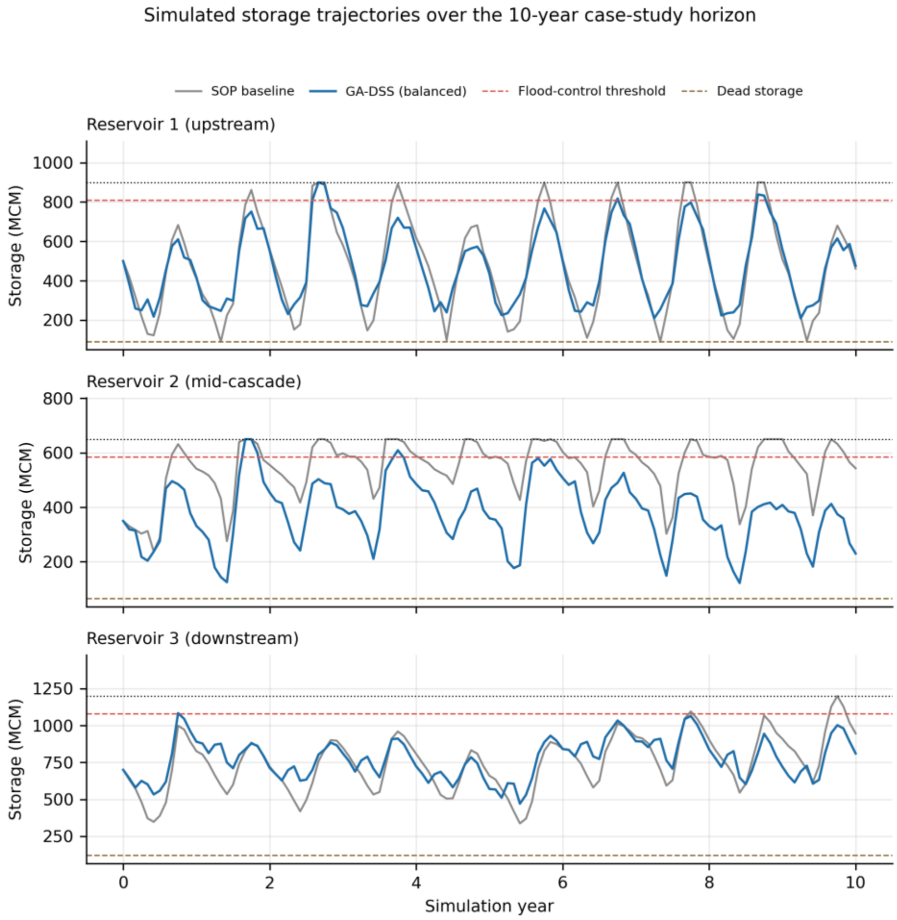

Figure 4. Simulated monthly storage trajectories for the GA-DSS balanced policy versus the Standard Operating Policy baseline, over the ten-year case-study horizon, at each of the three cascade reservoirs.

Reliability, Resilience and Vulnerability Analysis

Table 5 reports the Hashimoto et al. [12] reliability-resilience-vulnerability (RRV) indices computed independently at each reservoir node for the GA-DSS balanced policy and the SOP baseline. At Reservoir 2 and Reservoir 3, the SOP baseline achieves perfect reliability, resilience and vulnerability scores (no shortfall ever occurs at these nodes under SOP operation), whereas the GA-DSS policy introduces moderate, controlled shortfalls at these same nodes (reliability of 80.8% and 81.7%, respectively). This is the clearest evidence in this study of the classical hedging trade-off: the GA has learned to redistribute rationing across the entire cascade -- including nodes that would otherwise never experience shortage under greedy, demand-matching operation -- in order to retain system-wide storage for flood control and enhanced hydropower generation. At Reservoir 1, where water is already relatively scarce (the smallest capacity-to-mean-inflow ratio in the cascade), both policies experience some shortfall, and the SOP baseline in fact achieves higher reliability (81.7% versus 64.2%), but with comparable resilience (45.5% versus 46.5%) and vulnerability (71.7% versus 70.9%) -- meaning that although the GA-DSS policy fails somewhat more often at this specific node, its failures are, if anything, marginally less severe on average and recover at a statistically similar rate.

Table 5. Reliability-resilience-vulnerability (RRV) performance indices [12] by reservoir node: GA-DSS balanced policy versus SOP baseline.

Reservoir node	Operating policy	Reliability (%)	Resilience (%)	Vulnerability (% of demand)
Reservoir 1 (upstream)	GA-DSS (balanced)	64.2	46.5	70.9
Reservoir 1 (upstream)	SOP	81.7	45.5	71.7
Reservoir 2 (mid-cascade)	GA-DSS (balanced)	80.8	43.5	38.4
Reservoir 2 (mid-cascade)	SOP	100.0	100.0	0.0
Reservoir 3 (downstream)	GA-DSS (balanced)	81.7	50.0	62.2
Reservoir 3 (downstream)	SOP	100.0	100.0	0.0

Taken together, Sections 4.3-4.5 show that the GA-based decision support system does not uniformly dominate the SOP baseline -- nor should it be expected to, given that reservoir operation is a genuinely multi-objective problem in which water-supply reliability, hydropower generation and flood-control safety cannot all be maximised simultaneously. Instead, the DSS makes the shape of that trade-off explicit and quantified, and allows the reliability-vulnerability balance itself to become a managerial choice: Section 5 shows how a manager who places a higher priority on water-supply reliability can select a different point on the same Pareto front (the water-supply-priority recommendation in Table 3) that still improves flood-control performance by 78% relative to the SOP baseline while achieving higher hydropower generation, at a smaller reliability cost than the balanced recommendation.

THE GA-BASED MANAGERIAL DECISION SUPPORT SYSTEM

A Pareto-optimal front, however well populated, is not by itself an operational recommendation: a water utility or basin authority must still select a single operating policy to implement. The role of the decision support system developed in this study is to close that gap by combining the NSGA-II optimisation engine (Section 3.4) with a TOPSIS-based compromise-selection layer that converts a manager's stated preferences directly into a specific, recommended operating policy drawn from the consolidated Pareto front.

From Pareto Front to Operating Recommendation

Table 3 already reported the outcome of this process for three illustrative preference profiles. Under a balanced profile (equal weight on water-supply deficit, hydropower generation and flood control), the DSS recommends a policy achieving 6,564.2 GWh of ten-year hydropower generation with a flood-exceedance index of 4.13. Under a water-supply-priority profile (weight 0.6 on the deficit objective), the recommended policy shifts toward higher reliability (77.8% versus 75.6% for the balanced policy) while retaining strong flood-control performance (f3 = 9.00). Under a hydropower-priority profile (weight 0.6 on the energy objective), the DSS instead recommends the policy achieving the highest energy generation of the three profiles (6,810.0 GWh). Because all three recommendations are drawn from the same, once-computed Pareto front, re-running the compromise-selection step for a new preference profile is computationally instantaneous -- it involves no re-optimisation, only a weighted ranking of the already-known 417 candidate policies -- which is what makes this architecture practical as an interactive decision support tool rather than a one-shot analysis.

Decision Support Workflow and Deployment

In operational use, the workflow set out in this study is intended to run as follows. First, the NSGA-II optimisation (Section 3.4) is run once, offline, against the best available inflow forcing for the basin -- historical records from a network such as GRDC or USGS NWIS [17,18], a stochastically generated ensemble as used in this case study, or, where available, seasonal or climate-scenario forecasts -- to produce (or periodically refresh) the consolidated Pareto front. This step, requiring on the order of a few minutes of computation for a system of this size on a standard workstation, is the only computationally demanding part of the workflow and does not need to be repeated for every decision. Second, water resource managers interact with the resulting front through the TOPSIS compromise-selection layer, adjusting preference weights (or, equivalently, directly inspecting the Pareto front visualised as in Figure 2) to explore how different stated priorities translate into different operating recommendations, storage trajectories, and expected reliability, energy and flood-safety outcomes. Third, the recommended policy -- expressed as a small table of twelve monthly target releases plus three hedging/flood parameters per reservoir (45 numbers in total for the present three-reservoir system) -- can be handed directly to operations staff as a standing rule curve, audited against institutional or regulatory constraints, and updated on a rolling basis (for example, seasonally or annually) as new inflow data becomes available. Because every component of the optimiser is implemented transparently in this study rather than treated as a black-box library call, a utility's engineering staff can inspect, and if necessary adapt, every step of the recommendation process, addressing the auditability and trust concerns that are frequently cited as barriers to the operational adoption of optimisation-based reservoir management tools [3,4].

Limitations

Several limitations should guide the interpretation and any operational extension of this study. First, the case-study hydrology is a statistically representative synthetic sequence rather than a historical or forecast record for a specific named basin; before operational deployment, the same framework should be re-run against the actual gauged or forecast inflow record for the basin of interest, which the modelling architecture (Sections 3.1-3.2) supports without modification. Second, the operating policy is a monthly-time-step, storage-feedback hedging rule; basins requiring finer temporal resolution (e.g. daily flood routing) or explicit look-ahead based on seasonal forecasts would require extending the state representation, which would increase the decision-variable count and computational cost accordingly. Third, the three managerial preference profiles used to illustrate the TOPSIS layer are illustrative rather than elicited from an actual basin authority; operational use would involve a structured preference-elicitation exercise with the relevant stakeholders. Fourth, as with any evolutionary optimisation, the consolidated Pareto front is an approximation of the true Pareto-optimal set rather than a certified global optimum; the multi-replicate design and hypervolume convergence tracking used in this study (Section 4.1) mitigate, but cannot entirely eliminate, this approximation.

CONCLUSION

This study developed a complete, from-first-principles genetic-algorithm-based managerial decision support system for multi-objective reservoir operation and applied it to a three-reservoir cascade case study representative of a monsoon-influenced river basin. A real-coded, constraint-handling NSGA-II optimiser was coupled to a physically explicit, vectorised reservoir simulation model and used to search a compact, 45-parameter, cyclic hedging-rule operating-policy space -- chosen specifically so that the optimisation output is a genuinely repeatable, implementable operating rule rather than a one-off release schedule. Across eight independent replicates, the optimiser converged to a well-populated, 417-member consolidated Pareto front trading off water-supply deficit, hydropower generation and flood-control performance, with convergence confirmed by hypervolume tracking and robustness confirmed by a systematic parameter sensitivity analysis.

Relative to a conventional Standard Operating Policy expressed within the same policy structure, the three GA-DSS recommendations examined in this study increased hydropower generation by 1.7-5.5%, cut the flood-control exceedance index by 78-90%, and reduced spill by 78-81%, while deliberately trading away some period-wise water-supply reliability -- a reliability-resilience-vulnerability trade-off shown to be consistent with established hedging theory rather than an artefact of the optimisation, and one that a manager can tune directly through the TOPSIS-based compromise-selection layer. This layer is what elevates the optimisation from an academic Pareto-front exercise to a genuine decision support system: it lets a water resource manager translate stated priorities among water supply, energy and flood safety directly into a concrete, implementable operating recommendation, with no further optimisation required.

The principal practical advantage of the framework is transparency: every genetic operator, every simulation equation, and every metric reported in this study is implemented explicitly and reproducibly, so that the system can be audited, adapted and trusted by the engineering staff who would operate it, rather than treated as an opaque external tool. Its principal limitations are that the case-study hydrology is synthetic rather than site-specific, that the operating policy is resolved at a monthly time step, and that the illustrative preference profiles used to demonstrate the decision support layer should be replaced with weights elicited from the actual responsible authority before operational use. None of these limitations is structural: the same simulation, optimisation and decision-support architecture applies directly to a real basin once its historical or forecast inflow record is substituted for the synthetic sequence used here, making this framework a practical, extensible template for evidence-based, multi-objective reservoir management.

ACKNOWLEDGMENT

The authors gratefully acknowledge the constructive comments of the anonymous reviewers and editors, which improved the clarity and rigour of this manuscript.

Funding

The study was funded by the Foundation for Scientific and Technological Development of Yugra under project number № 2026-252-01.

CONFLICT OF INTERESTS

The authors declare that there is no conflict of interest associated with this publication.

NOMENCLATURE

Table N1. Principal symbols used in this study.

Symbol	Definition
$S_{i,t}$	Storage of reservoir i at the start of period t (MCM)
$I_{i,t}$	Inflow to reservoir i in period t (MCM/month)
$E_{i,t}$	Evaporation loss from reservoir i in period t (MCM/month)
$R_{i,t}$	Decision (and actual) release from reservoir i in period t (MCM/month)
$Sp_{i,t}$	Spill from reservoir i in period t (MCM/month)
$T_{i,m}$	GA-optimised target release of reservoir i for calendar month m (MCM/month)
h_i, k_i, g_i	Hedging-trigger fraction, hedging-severity exponent, and flood-release gain of reservoir i

$D_{i,m}$	Consumptive (irrigation + municipal) demand at node i in calendar month m (MCM/month)
$f_{div,i}$	Consumptive diversion fraction of release at node i
$H_{i,t}$	Net hydraulic head at reservoir i in period t (m)
η_i	Turbine-generator efficiency of reservoir i
f1, f2, f3	Water-supply deficit, (negative) hydropower energy, and flood-exceedance objectives
g1	Aggregate end-of-horizon carry-over storage constraint violation
w	Vector of managerial preference weights used in TOPSIS compromise selection

REFERENCES

1. Holland, J.H. *Adaptation in Natural and Artificial Systems*; University of Michigan Press: Ann Arbor, MI, USA, 1975.
2. Goldberg, D.E. *Genetic Algorithms in Search, Optimization, and Machine Learning*; Addison-Wesley: Reading, MA, USA, 1989.
3. Nicklow, J.; Reed, P.; Savic, D.; Dessalegne, T.; Harrell, L.; Chan-Hilton, A.; Karamouz, M.; Minsker, B.; Ostfeld, A.; Singh, A.; et al. State of the art for genetic algorithms and beyond in water resources planning and management. *J. Water Resour. Plann. Manag.* **2010**, *136*, 412–432. https://doi.org/10.1061/(ASCE)WR.1943-5452.0000053.
4. Maier, H.R.; Kapelan, Z.; Kasprzyk, J.; Kollat, J.; Matott, L.S.; Cunha, M.C.; Dandy, G.C.; Gibbs, M.S.; Keedwell, E.; Marchi, A.; et al. Evolutionary algorithms and other metaheuristics in water resources: Current status, research challenges and future directions. *Environ. Model. Softw.* **2014**, *62*, 271–299. https://doi.org/10.1016/j.envsoft.2014.09.013.
5. Labadie, J.W. Optimal operation of multireservoir systems: State-of-the-art review. *J. Water Resour. Plann. Manag.* **2004**, *130*, 93–111. https://doi.org/10.1061/(ASCE)0733-9496(2004)130:2(93).
6. Wardlaw, R.; Sharif, M. Evaluation of genetic algorithms for optimal reservoir system operation. *J. Water Resour. Plann. Manag.* **1999**, *125*, 25–33. https://doi.org/10.1061/(ASCE)0733-9496(1999)125:1(25).
7. Chen, L. Real coded genetic algorithm optimization of long term reservoir operation. *JAWRA J. Am. Water Resour. Assoc.* **2003**, *39*, 1157–1165. https://doi.org/10.1111/j.1752-1688.2003.tb03699.x.
8. Reddy, M.J.; Kumar, D.N. Multi-objective particle swarm optimization for generating optimal trade-offs in reservoir operation. *Hydrol. Process.* **2007**, *21*, 2897–2909. https://doi.org/10.1002/hyp.6507.
9. Deb, K.; Pratap, A.; Agarwal, S.; Meyarivan, T. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans. Evol. Comput.* **2002**, *6*, 182–197. https://doi.org/10.1109/4235.996017.
10. Deb, K.; Agrawal, R.B. Simulated binary crossover for continuous search space. *Complex Syst.* **1995**, *9*, 115–148.

11. Deb, K. An efficient constraint handling method for genetic algorithms. *Comput. Methods Appl. Mech. Eng.* **2000**, *186*, 311–338. https://doi.org/10.1016/S0045-7825(99)00389-8.
12. Hashimoto, T.; Stedinger, J.R.; Loucks, D.P. Reliability, resiliency, and vulnerability criteria for water resource system performance evaluation. *Water Resour. Res.* **1982**, *18*, 14–20. https://doi.org/10.1029/WR018i001p00014.
13. Hwang, C.L.; Yoon, K. *Multiple Attribute Decision Making: Methods and Applications*; Springer-Verlag: Berlin/Heidelberg, Germany, 1981. https://doi.org/10.1007/978-3-642-48318-9.
14. Thomas, H.A.; Fiering, M.B. Mathematical synthesis of streamflow sequences for the analysis of river basins by simulation. In *Design of Water Resource Systems*; Maass, A., Hufschmidt, M.M., Dorfman, R., Thomas, H.A., Marglin, S.A., Fair, G.M., Eds.; Harvard University Press: Cambridge, MA, USA, 1962; pp. 459–493.
15. Loucks, D.P.; van Beek, E. *Water Resource Systems Planning and Management: An Introduction to Methods, Models, and Applications*, 2nd ed.; Springer: Cham, Switzerland, 2017. https://doi.org/10.1007/978-3-319-44234-1.
16. Zitzler, E.; Thiele, L. Multiobjective evolutionary algorithms: A comparative case study and the strength Pareto approach. *IEEE Trans. Evol. Comput.* **1999**, *3*, 257–271. https://doi.org/10.1109/4235.797969.
17. U.S. Geological Survey (USGS). National Water Information System: Web Interface. Available online: https://waterdata.usgs.gov/nwis (accessed on 12 Jan 2026).
18. Global Runoff Data Centre (GRDC). The GRDC Data Portal. Available online: https://grdc.bafg.de (accessed on 12 Jan 2026).