Optimality

www.opt.reapress.com

Opt. Vol. 3, No. 3 (2026) 234–247.

Paper Type: Original Article

Adaptive Evolutionary Search in Changing Environments: Moving Peaks and a Two-Species Chemostat Surrogate

Seyyed Farid Ghannadpour[1,*], **Daniel A. Ferreira**[2]

[1] Department of Industrial Engineering, Iran University of Science and Technology, Tehran, Iran; ghannadpour@iust.ac.ir.
[2] Department of Electrical and Computer Engineering, University of Campinas, Campinas, Brazil; daniel.ferreira@unicamp.br.

Citation:

Received: 12 March 2026 *Revised: 26 May 2026* *Accepted: 19 July 2026*	Ghannadpour, S. F., & Ferreira, D. A. (2026). Adaptive evolutionary search in changing environments: Moving peaks and a two-species chemostat surrogate. *Optimality, 3*(3), 234-247.

Abstract

Dynamic optimization can support the search for operating designs when a simulated nutrient regime changes during an optimization run. This study reports AMPERE, a memory-based evolutionary search combining sentinel monitoring, a regime-indexed archive, empirical transition estimates, allocation of search effort and a magnitude-scaled mutation response. The reported campaign compares six methods on a moving-peaks landscape with 62,500 evaluations and a fixed two-species chemostat surrogate with 17,500 evaluations, using ten main-comparison seeds and five ablation seeds. On moving peaks, AMPERE has a larger within-epoch incumbent-drop metric than the plain GA (0.293 ± 0.124 versus 0.103 ± 0.031; Holm $p = 0.0098$), but a smaller value than fixed triggered hypermutation (2.622 ± 0.788). On the chemostat surrogate, its mean current-best endpoint titer is 4.08 ± 0.17 g/L, compared with 1.55 ± 0.41 g/L for the plain GA. Its comparisons with random immigrants and triggered hypermutation do not reject the null hypothesis of equal performance at the reported significance level. Removing AMPERE's mutation response gives 1.76 ± 0.50 g/L in the reported ablation; the five-pair two-sided p-value is 0.0625. These observations show that performance depends on the benchmark, without establishing an isolated benefit of anticipation or memory. The incumbent-drop metric is not regret relative to a known optimum, and neither benchmark validates biological adaptation or online plant control. Exact simulator parameters, paired outputs, event counters and implementation details remain necessary for independent reproduction.

Keywords: Dynamic optimization, Evolutionary algorithms, Stochastic environments, Microbial communities, Chemostat, Change detection.

1 | Introduction

Quantitative microbial-growth models and the chemostat provide a basis for studying nutrient-limited cultures [1–3]. Laboratory evolution has also characterized adaptation and diversification across generations [4], [5], while stochastic-switching models connect phenotypic diversity and environmental information [6],

Corresponding Author: ghannadpour@iust.ac.ir
https://doi.org/10.22105/opt.v3i3.117
Licensee System Analytics. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (http://creativecommons.org/licenses/by/4.0).

[7]. These studies motivate fluctuating-environment problems without establishing that candidate changes in an optimizer reproduce microbial evolution.

Bioreactor operation can involve both uncertain parameters and changing feed conditions [8], [9]. Community modeling adds interactions among organisms [10], [11]. A search algorithm can therefore be evaluated against changing simulated fitness functions while remaining distinct from the physiological or evolutionary adaptation of a real community. The present study concerns the former: designs are optimizer candidates, and their performance is produced by a simulator.

Evolutionary dynamic optimization uses memory, diversification and multiple populations to respond to changing objective functions [12–18]. Memory can reuse stored designs, while immigrants and triggered variation provide new candidates. Predictive allocation adds another assumption: estimated environmental transitions must carry information useful enough to justify diverted evaluations.

This study brings established dynamic-search mechanisms to a particular two-species chemostat surrogate. Its contribution is a bounded comparison of a combined search design with simpler baselines on two reported benchmark configurations. It does not establish that dynamic optimization and bioprocess modeling were previously disjoint or that the proposed modules are individually novel. Abstract landscapes and mechanistic simulators supply complementary tests, but a difference in their rankings is not itself evidence of transfer to industrial communities.

The formulation distinguishes optimizer evaluations from physical simulation time and distinguishes a per-regime optimum from the within-method reference used in the reported metric. AMPERE combines established detection, memory, allocation and mutation ideas. The comparison asks whether this combination improves the reported objective and recovery summaries at the stated evaluation budgets.

Section 2 relates the search design to dynamic optimization and bioprocess modeling. Section 3 states the questions and separates the theoretical objective from measured quantities. Section 4 describes the reported AMPERE modules and their implementation dependencies. Section 5 interprets the original numerical summaries, and Sections 6 and 7 discuss limitations and further tests.

2 | Related Work

2.1 | Evolutionary Dynamic Optimization

Early work examined diploid representations and diversity in changing environments [12], [14], [15]. Later memory and multi-population methods stored useful solutions or maintained search in multiple regions [13], [16], [21]. These sources establish antecedents for AMPERE's modules; they do not establish that the particular combined implementation is superior.

Dynamic test generators and moving-peaks configurations provide controlled ways to study changing fitness landscapes [13], [19], [20]. Surveys distinguish memory, diversity and prediction mechanisms [17], [22]. The reported campaign uses measures of current-best fitness and recovery together with a within-method incumbent-drop measure. The latter differs from benchmark error computed relative to a known current optimum, so it should not be read as an interchangeable measure of tracking accuracy.

Operator adaptation and uncertainty handling are also covered in the dynamic-optimization literature [13], [18]. Their relevance depends on the change process and evaluation constraints. The present transition estimate represents only the reported small regime space and does not demonstrate accurate prediction across broader environment classes.

2.2 | Microbial Dynamics in Variable Environments

Chemostat growth and inhibition models provide one basis for interpreting nutrient-dependent dynamics [1–3], [23]. Classical population models and experiments [24], [25] and later microbial-interaction studies [11], [26–28] concern ecological structure and stability. The present experiment uses a fixed two-species surrogate

with continuous candidate variables; it does not optimize community membership or certify ecological stability. Genome-scale and dynamic flux-balance methods [29], [30] remain related modeling approaches rather than executed components.

Phenotypic switching, bet-hedging and persistence provide biological reasons to study fluctuating environments [6], [7], [31]. These mechanisms concern organisms and populations. They motivate the application setting, but do not validate the numerical response of an evolutionary search algorithm or establish that simulated design changes correspond to realizable strain adaptation. Here the species set is fixed and a stylized feed-matching trait is an optimizer variable.

2.3 | Bioprocess Optimization and Control under Variability

Bioprocess estimation and control address changing states and uncertain kinetics [8], [9]. Robust optimization also covers multiple uncertainty settings and cannot be restricted categorically to a static environment [32]. These approaches are plausible comparators, but the present campaign includes neither a robust-design optimizer nor model-predictive control. Metabolic-community and dynamic flux-balance models offer richer biological representations [29], [30]. However, the reported chemostat surrogate does not execute a genome-scale flux solve. ANN–metaheuristic optimization has also been applied to enzyme and biomass responses as functions of fermentation-medium composition [44]. That static single-organism response problem differs from the changing two-species dynamical surrogate considered here.

2.4 | Positioning of the Comparison

The relevant research question is how search mechanisms behave when fitness is generated by a changing simulator. *Table 1* summarizes potential roles of three research traditions rather than asserting that any tradition lacks an entire capability. The numerical comparison is restricted to the configured methods and models. It neither demonstrates an exhaustive literature gap nor measures an industrial advantage.

Table 1. Complementary roles of research traditions; entries describe typical contributions rather than categorical exclusions.

Search Role	Evolutionary Dynamic Optimization	Microbial Dynamics Modelling	Bioprocess Control
Candidate exploration	Population and trajectory search	Model-dependent design studies	Optimization can support control
Biological representation	Can use a mechanistic simulator	Interactions and growth mechanisms	Model-dependent process dynamics
Changing environments	Detection, memory and diversification	Variable-environment analysis	Estimation and disturbance handling
Prediction	Model-dependent anticipatory search	Ecological and evolutionary models	Predictive control where specified
Stored information	Solution archives where used	Model and state histories	Estimated states and controller memory

3 | Problem Formulation and Research Questions

3.1 | Research Questions

Three questions organize the comparison. The first asks whether regime-indexed memory improves recovery when simulated regimes recur; the second asks whether transition-weighted allocation improves performance beyond reactive diversification; the third asks whether the ranking of the configured methods differs between the two benchmarks. The campaign can describe these configured outcomes, but does not vary recurrence, contrast and model class independently.

The working expectation is that reusable elites may shorten recovery when a relevant regime returns. A matched-budget comparison with memory disabled would be needed to isolate this effect; a lack of statistical rejection would not establish equivalence.

The allocation question requires distinguishing predictions of successor regimes from fitness calls used to prepare candidates for them. Every such call, including a counterfactual-regime simulation, must be included in the evaluation accounting.

The ranking question concerns the particular moving-peaks and chemostat instances. Differences between them can arise from geometry, contrast, objective scale, recurrence, budget and operator settings; this comparison does not identify a unique cause.

3.2 | System Model and Notation

The executed nutrient process is described as a discrete two-state Markov chain [33] at scheduled optimizer epoch boundaries, with a separate fine-step feed chain inside each 96-hour simulation. It is not a documented continuous-time jump-process implementation. The species set is fixed at two; the exact continuous design coordinates, bounds, initial states and kinetic constants must be recovered from the original simulator. *Table 2* separates physical-time quantities from optimizer-indexed quantities.

Table 2. Principal notation for the dynamic community optimization problem.

Symbol	Meaning	Unit
z_e	macro nutrient regime at optimizer epoch e	dimensionless
Z	finite regime set	dimensionless
P	discrete macro-regime transition matrix	dimensionless
x_i(t)	biomass concentration of species i	g L^{-1}
s_k(t)	substrate concentration k	g L^{-1}
μ_i(s; z)	specific growth rate	h^{-1}
θ	continuous candidate design; species set fixed	coordinate-dependent
f(θ; z)	96-hour endpoint titer in the chemostat	g L^{-1}
E_off	within-method incumbent drop averaged over optimizer steps	fitness units
τ_rec	censored first crossing of the pooled epoch reference	optimizer generations

3.3 | Formal Problem Statement

Eq. (1) gives a generic biomass balance for interpreting a chemostat model, rather than a complete specification of the executed simulator. A full model also requires substrate, by-product and product balances, uptake functions, yields, initial states and the decision-to-parameter map. *Eq. (2)* defines the reported fitness as endpoint product concentration. *Eq. (3)* states the discrete macro-regime transition model; a symmetric switching probability of 0.6 corresponds to off-diagonal entries of 0.6 and diagonal entries of 0.4. Here x_i is biomass concentration, D is dilution rate and k_d,i is a loss or decay rate, with both rates in h^{-1}. P_{θ,z}(t) denotes product concentration in g/L for decision θ and regime z; this scalar concentration is distinct from the transition matrix P in *Eq. (3)*.

$$\frac{dx_i}{dt} = [\mu_i(s; z) - D - k_{d,i}]x_i. \quad (1)$$

$$f(\theta; z) = P_{\theta,z}(96\,h). \quad (2)$$

$$\Pr(z_{e+1} = j \mid z_e = i) = P_{ij}, \quad P_{ij} \geq 0, \quad \sum_j P_{ij} = 1. \quad (3)$$

For a theoretical tracking objective, let π be a causal search policy producing a design θg at optimizer index g, with access only to observations available by that index. Let f*(z) be the supremum of the expected endpoint fitness over the stated design box in regime z. *Eq. (4)* gives an idealized expected gap for this policy. This comparator is not computed by the reported campaign, whose finite probes do not certify f*(z).

$$\min_{\pi} \mathrm{E}[\frac{1}{G}\sum_{g=1}^{G}(f^{*}(z_g) - f(\theta_g^{\pi}; z_g))]. \quad (4)$$

Eq. (4) optimizes a policy rather than a single fixed design. It contains a fitness-gap term and no switching-cost term. Simulation time inside a fitness evaluation and optimizer index across evaluations serve different roles. The reported mean current-best titer averages optimizer outcomes from reset simulation windows; it is not time-averaged production from one continuously operated reactor.

4 | Proposed Methodology: AMPERE

AMPERE denotes the proposed combination of sentinel monitoring, regime-indexed elites, empirical transition estimates, transition-weighted allocation and a mutation response. The following description is restricted to mechanisms reported for the numerical configuration. More elaborate classification, sensing, forgetting and indexing policies require the original implementation and are not assumed to have been executed.

4.1 | Environment Monitoring and Change Detection

The reported detector re-evaluates five fixed sentinel designs each generation and combines a two-sided 5 percent fitness-displacement rule with cumulative-sum monitoring [34]. The supplied description does not resolve the CUSUM threshold, warm-up, reset policy or whether the two rules are joined by conjunction, disjunction or a confidence-dependent gate. Detection results are therefore reported without attributing missed events to a particular rule.

The observations in this campaign are simulator fitness values. Online substrate, off-gas or product measurements from a biological process were not tested. The deterministic evaluation setting does not establish robustness to measurement noise or quantify the separate benefit of cumulative-sum monitoring.

4.2 | Dual Archive: Regime Memory and Environment Model

The reported memory contains regime-indexed elites and a recent-elite fallback. Stored designs are re-evaluated before reuse. The transition estimate uses observed regime transitions with a Dirichlet prior. Exact elite capacity, regime-assignment rule, count update and reset conventions are implementation dependencies; nearest-centroid classification and fitted dwell-time distributions are not documented executed components.

Forgetting may be useful when stored solutions become stale, but exponential discounting of transition counts does not by itself bound the number of archived regimes or designs. A finite capacity and deletion policy would be required for a memory bound. Such a policy should be specified from the source code before any bounded-memory claim is made.

4.3 | Anticipatory Allocation of Search Effort

The reported allocation assigns search effort using estimated successor probabilities and activates at a successor probability of 0.25. Population sizes, floors, migration and the handling of counterfactual-regime evaluations need an explicit implementation record. Cached regime parameters do not make new fitness evaluations free; their simulation calls must be charged to the same budget.

Allocation can be viewed as a choice between current tracking and preparation for possible future regimes. Its benefit depends on prediction quality and the cost of diverted evaluations. No-free-lunch results [35] do not prove a benefit or a failure on either restricted benchmark family; those questions require the relevant controlled comparisons.

4.4 | Change Response and Main Loop

The reported mutation response increases with detected change magnitude and is capped when re-evaluated elites remain competitive. Its exact scaling function, restart fraction and cap require source-code verification. The fixed-hypermutation comparator multiplies the mutation scale by ten for ten generations, providing a distinct configured response rather than an otherwise identical controller.

The executed GA uses tournament selection of size three, simulated binary crossover with index 15 and polynomial mutation with index 20, without elitism. There is no reported strain-membership bit string in this fixed two-species experiment. The operational sequence below is a summary of the reported modules, not a replacement executable specification.

The search begins with the candidate population, sentinel designs and empty or initialized memory records. Within an epoch, it evaluates and varies designs under the active fitness configuration while updating sentinel statistics. Upon a detected change it updates available regime records, re-evaluates reusable elites, applies the configured mutation response and updates allocation. Execution stops at the shared evaluation budget; implementation-dependent ordering, accounting and resets require the original code.

4.5 | Complexity, Assumptions, and Failure Conditions

For population size N and s sentinel designs, with a denoting the number of archive re-evaluations at each change, the immediate evaluation overhead is s calls per monitored generation and a calls per archive refresh. Total runtime also depends on simulator cost and counterfactual evaluations. No logarithmic archive retrieval or bounded-growth guarantee follows from the supplied description. Matched total evaluation limits can constrain calls without making offspring counts or generation counts identical.

Memory and anticipation require reusable regime structure and informative observations. Noise, misclassification and changing kinetics can degrade both modules. The reported campaign does not independently manipulate these factors, so it cannot identify a recurrence threshold or a noise level at which AMPERE becomes preferable.

False-negative and false-positive detection rates require an event-level record with a defined denominator, tolerance window and treatment of repeated alarms. Reported detection totals alone do not provide both rates.

Unbounded or continuously changing regime spaces also require a specified clustering and archive-capacity policy. Simulator–process mismatch remains a separate external-validity issue and was not measured in either benchmark.

5 | Experimental Results and Analysis

5.1 | Setup: Benchmarks, Methods, Protocol, and Scope

The reported comparison uses six population-based methods. Moving peaks has five dimensions on $[0,100]^5$, ten sphere peaks with heights initially uniform on [30,70] and widths on [1], [5], shift severity 1.0 and a clipped Gaussian height walk of severity 7.0. Changes occur every 2,500 evaluations; each run has 25 epochs and 62,500 evaluations, with 24 scheduled changes and population size 50. The reported configuration is non-recurrent; its exact generator and peak-update conventions remain necessary for reproduction [13], [19].

The chemostat surrogate fixes a producer and a consumer. The producer forms an inhibitory by-product and product; the consumer removes by-product while incurring maintenance and product-consumption costs. A stylized producer trait has different preferred values under the two macro regimes. The reported macro-state chain switches with probability 0.6 at scheduled boundaries. Inside each reset 96-hour simulation, feed switching is approximately 0.001 per 0.5-hour step. This corresponds to about 0.002 switches per hour, or a discrete-step hazard $-\log(0.999)/0.5 \approx 0.002001\ h^{-1}$, rather than a continuous-time intensity of $0.001\ h^{-1}$.

Explicit Euler uses 192 steps with positivity clipping. Fitness is endpoint titer in g/L. Seven optimizer epochs give 17,500 evaluations and six scheduled boundaries per seed; a boundary need not produce an actual macro-state switch.

The reported baselines are a generational GA with tournament size three, crossover index 15, mutation index 20 and no elitism; random immigrants replacing 20 percent of offspring [14], [15], memory-enhanced search re-evaluating ten recent epoch elites and reinserting the five best [13], [36], triggered hypermutation multiplying mutation scale tenfold for ten generations [15], and five subpopulations of ten with ring exchange every five generations [16]. AMPERE adds the reported regime archive, Dirichlet-regularized transition estimate, allocation threshold of 0.25 and magnitude-scaled response. Detection-equipped methods share the stated five-sentinel detector, but other module and accounting differences remain relevant to attribution.

The source reports five sentinel calls per monitored generation, ten archive re-evaluations per change, common environmental realizations within each main-comparison seed, ten main seeds, and five ablation seeds. Exact counters are needed to verify that all calls, including successor-regime tests, fall within the budget. NumPy [37], SciPy [38] and Matplotlib [39] are reported, with a total runtime of 304 seconds on one core; the original run log is needed to attest that measurement. A 400-point Latin-hypercube probe produced reported best titers of 4.09, 4.08, 4.07, and 3.99 g/L, as well as trait values of 1.29 and 0.79 in the contrasted configurations. These are finite-sample observations, not proven regime optima or ceilings.

For method m and epoch e, let bmg be its current-best fitness at optimizer step g and hmg = max{bmu: u is in the same epoch and u ≤ g}. The reported offline-error column averages hmg − bmg. It therefore measures a loss relative to that method's own running incumbent, not to a known optimum, and can remain zero for consistently poor designs. The final-window error uses the last five generations. Recovery is the first generation reaching 95 percent of the best value attained by any method in the same seed–epoch pool; non-attainment is recorded at the epoch length. Such a censored, retrospectively pooled measure is not an uncensored expected hitting time. Because overhead changes the number of generations, evaluation-index recovery would be needed for a stricter budget comparison.

5.2|Moving-Peaks Comparison

Table 3 reports the moving-peaks summaries, and Fig. 1 shows a representative aggregate trajectory. The plain GA has the smallest incumbent-drop metric (0.103 ± 0.031), followed by memory-enhanced and multi-population GAs (0.125 ± 0.056 and 0.156 ± 0.055). AMPERE gives 0.293 ± 0.124, below random immigrants (0.557 ± 0.415) and fixed hypermutation (2.622 ± 0.788). These differences are compatible with disruptive variation under small shifts, but aggregate traces do not isolate that mechanism. Multi-population search has the largest reported mean current-best fitness, 53.1 ± 8.7, and the fastest pooled-reference recovery, 24.1 ± 8.0 generations. Thus the incumbent-drop ordering is not the complete objective-quality ordering. The detector reports 219 detections among 240 scheduled changes, corresponding to 91.25%; event logs are required to establish missed-event causes or false-alarm rates.

Table 3. Moving-peaks results: Within-method incumbent-drop metric, censored pooled-reference recovery, final-window drop, mean current-best fitness and reported detections (mean ± standard deviation over ten seeds).

Method	Offline Error	Recovery Gens	Final-Window Error	Mean Current Best	Changes Detected
MPGA	0.156 ± 0.055	24.1 ± 8.0	0.256 ± 0.207	53.1 ± 8.7	-
GA	0.103 ± 0.031	33.1 ± 11.5	0.124 ± 0.069	47.0 ± 8.8	-
MEGA	0.125 ± 0.056	28.2 ± 10.1	0.211 ± 0.208	48.1 ± 10.1	219/240
RIGA	0.557 ± 0.415	33.4 ± 11.2	2.082 ± 4.930	46.6 ± 8.8	-
AMPERE	0.293 ± 0.124	28.5 ± 10.4	0.116 ± 0.066	47.6 ± 8.1	219/240
THGA	2.622 ± 0.788	28.1 ± 10.8	0.241 ± 0.360	45.7 ± 8.6	219/240

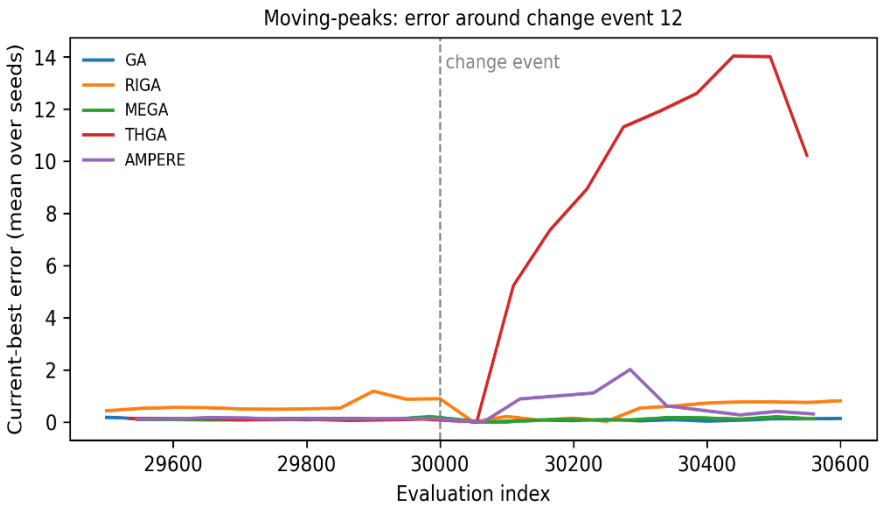


Fig. 1. Reported moving-peaks current-best-error trace around event 12, averaged over ten seeds. The plotted error must be interpreted with the source reference convention; the figure does not establish an isolated causal effect of a module.

5.3 | Component Ablations

Table 4 reports component ablations. On moving peaks, disabling the mutation response gives an incumbent-drop metric of 0.111 ± 0.075, compared with 0.293 ± 0.124 for the displayed full method. The reported paired W = 1.0 and p = 0.125 do not establish improvement at the 0.05 level or agreement in every nonzero pair. Removing the archive gives 0.424 ± 0.244. On the chemostat, disabling the response gives 1.76 ± 0.50 g/L rather than the displayed 4.08 ± 0.17 g/L, with W = 0 and p = 0.0625. For five nonzero paired differences under the exact sign-randomization null hypothesis, the smallest two-sided p-value is $2/2^5 = 0.0625$. Archive, allocation and transition ablations have displayed means 4.14 ± 0.02, 4.15 ± 0.02 and 4.15 ± 0.01 g/L, without a significant reported contrast. This does not establish equivalence, redundancy, or a uniquely causal module effect. The full rows duplicate main-comparison summaries; the matching five-seed full-method outputs must be supplied to verify the paired ablation statistics.

Table 4. Reported AMPERE ablations. MP denotes moving peaks and CH the chemostat. Ablation variants use five seeds; the provenance of the displayed full-method reference rows and matching paired outputs requires confirmation.

Variant	Offline Error (MP)	Mean Best (MP)	Offline Error (CH)	Mean Best Titer (CH)	Recovery Gens (CH)
AMPERE	0.293 ± 0.124	47.6 ± 8.1	0.047 ± 0.132	4.08 ± 0.17	2.1 ± 0.7
AMPERE - no - archive	0.424 ± 0.244	44.7 ± 8.0	0.006 ± 0.012	4.14 ± 0.02	1.9 ± 0.7
AMPERE - no - hyp	0.111 ± 0.075	48.6 ± 11.9	1.683 ± 0.428	1.76 ± 0.50	9.2 ± 5.4
AMPERE - no - alloc	not run	not run	0.000 ± 0.000	4.15 ± 0.02	1.5 ± 0.1
AMPERE - no - trans	not run	not run	0.000 ± 0.000	4.15 ± 0.01	1.4 ± 0.5

5.4 | Chemostat Comparison

Table 5 reports the chemostat summaries. AMPERE's mean current-best endpoint titer is 4.08 ± 0.17 g/L, while random immigrants and triggered hypermutation give 4.14 ± 0.01 and 4.12 ± 0.02 g/L. Plain, multi-population and memory-enhanced GAs give 1.55 ± 0.41, 1.55 ± 0.42 and 1.62 ± 0.41 g/L. The finite-probe value of 4.10 g/L is not a ceiling: two displayed method means exceed it. The source reports recovery

censoring at 36, 36 and 35 of 60 events for the latter methods, and detection totals of 40/40 for detector-equipped methods. These denominators cannot be identified with one another without an event log distinguishing scheduled boundaries, actual macro switches and eligible detection or recovery events. No failure mechanism is inferred from these totals alone.

The two benchmarks give different configured rankings, but differ in budget, geometry, fitness scale and regime construction as well as recurrence. The comparison therefore does not isolate recurrence or establish quantitative transfer. On the chemostat, low within-method drop values coexist with low titers, so inference uses titer rather than that drop metric. The reported comparisons do not isolate an additional benefit of transition prediction or regime-indexed memory beyond diversification.

Table 5. Chemostat results: Within-method incumbent-drop metric, censored pooled-reference recovery, mean current-best endpoint titer and reported detections (mean ± standard deviation over ten seeds). Event-denominator definitions remain an original-log dependency.

Method	Offline Error	Recovery Gens	Mean Best Titer (g/L)	Changes Detected
AMPERE	0.047 ± 0.132	2.1 ± 0.7	4.08 ± 0.17	40/40
RIGA	0.001 ± 0.002	1.9 ± 0.3	4.14 ± 0.01	-
THGA	0.000 ± 0.000	2.1 ± 0.4	4.12 ± 0.02	40/40
GA	0.000 ± 0.000	30.5 ± 4.2	1.55 ± 0.41	-
MPGA	0.000 ± 0.000	30.6 ± 4.3	1.55 ± 0.42	-
MEGA	0.000 ± 0.001	27.0 ± 3.8	1.62 ± 0.41	40/40

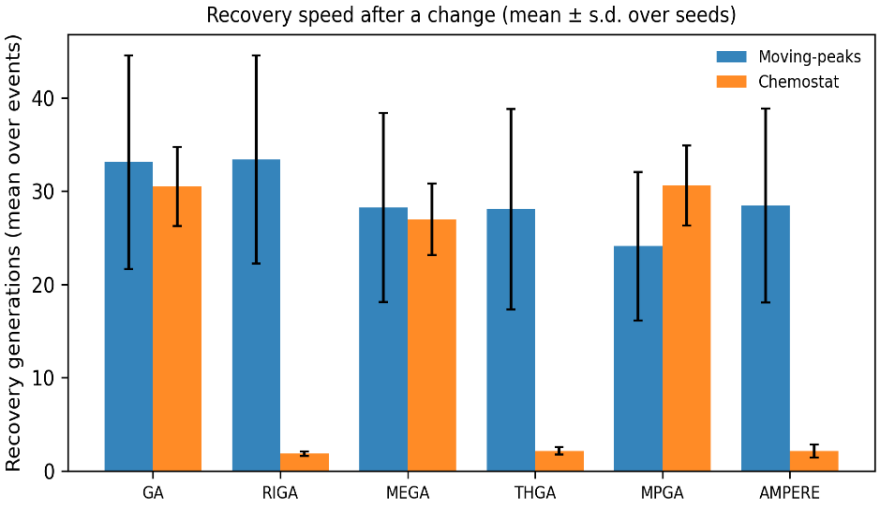


Fig. 2. Reported pooled-reference recovery summaries for the two benchmarks. Values include epoch-length censoring and use method-dependent generation clocks; they are not physical reactor recovery times.

5.5 | Statistical Analysis

Table 6 preserves the reported Friedman and paired signed-rank summaries [40–43]. The moving-peaks Friedman statistic 41.2 with five degrees of freedom has an asymptotic chi-square tail of approximately 8.55×10^{-8}; the displayed 0.0000 is rounded, not zero. AMPERE has larger incumbent-drop values than the plain and memory-enhanced GAs (Holm $p = 0.0098$ and 0.0176), but smaller values than triggered hypermutation (Holm $p = 0.0098$). Its contrast with multi-population search does not reject at 0.05 (Holm $p = 0.0547$), and it is numerically better than random immigrants with an unadjusted $p = 0.1934$. For the chemostat, the source reports chi-square = 40.7 and $p = 1.09 \times 10^{-7}$; using the rounded statistic gives approximately 1.08×10^{-7}. AMPERE has higher titer than the plain, memory-enhanced and multi-population GAs (Holm $p = 0.0098$ for each). The random-immigrant and triggered-hypermutation comparisons do not reject; this is not evidence of equivalence. Main-test pairing, tie handling, correction families and the five-seed ablation reference require

the original arrays. Seeds from a single parameterized instance do not represent ten independently sampled biological systems.

Table 6. Reported inferential summaries. A displayed p = 0.0000 denotes rounding. Direction is determined by the relevant metric and method summaries, not by the unsigned W statistic. Original paired arrays and correction-family definitions are required for verification.

Comparison	Statistic	p	Holm p
MP: Friedman, 6 methods (offline error)	41.2	0.0000	-
MP: AMPERE vs GA (offline error)	0.0	0.0020	0.0098
MP: AMPERE vs MEGA (offline error)	2.0	0.0059	0.0176
MP: AMPERE vs MPGA (offline error)	6.0	0.0273	0.0547
MP: AMPERE vs RIGA (offline error)	14.0	0.1934	0.1934
MP: AMPERE vs THGA (offline error)	0.0	0.0020	0.0098
CH: Friedman, 6 methods (mean best titer)	40.7	0.0000	-
CH: AMPERE vs RIGA (mean best titer)	17.0	0.3223	0.6445
CH: AMPERE vs THGA (mean best titer)	21.0	0.5566	0.6445
CH: AMPERE vs MEGA (mean best titer)	0.0	0.0020	0.0098
CH: AMPERE vs GA (mean best titer)	0.0	0.0020	0.0098
CH: AMPERE vs MPGA (mean best titer)	0.0	0.0020	0.0098
CH ablation: AMPERE vs no-no_archive (mean best titer)	7.0	1.0000	-
CH ablation: AMPERE vs no-no_alloc (mean best titer)	5.0	0.6250	-
CH ablation: AMPERE vs no-no_hyp (mean best titer)	0.0	0.0625	-
CH ablation: AMPERE vs no-no_trans (mean best titer)	4.0	0.4375	-

5.6 | Interpretation of the Reported Results

The reported pattern depends on the benchmark. AMPERE loses to the plain and memory-enhanced GAs on the moving-peaks incumbent-drop metric, while the multi-population GA leads mean current-best fitness. On the chemostat, AMPERE and two simpler diversification methods have higher reported mean endpoint titers than three less disruptive search methods. The mutation-response ablation is consistent with a useful response in this constructed chemostat, but its five-pair result is not significant at 0.05 and does not prove a unique mechanism. Memory, transition modeling and allocation remain unisolated rather than proven redundant.

The detector and archive consume fitness calls that could otherwise produce offspring. Their fraction depends on the realized generation and change counts; a universal 10 percent overhead is not established. The two fitness maps are reported as deterministic, so robustness to noisy biological observations remains untested. The two-species surrogate uses a deliberately contrasting trait and a coarse clipped Euler integrator. Integration convergence, parameter sensitivity and validation against measurements are needed before biological or process-level conclusions can be made.

These observations motivate a controlled follow-up that varies regime recurrence and contrast independently, uses matched evaluation accounting, and includes separate ablations for detection, memory, prediction, and mutation. They do not identify a general operating rule for an industrial bioprocess. A practically useful controller would additionally require representation of feasible actions, sensing delays and costs, state continuity, uncertainty, and model–process mismatch.

The numerical record consists of reported manuscript summaries and embedded figures. The original experiment program, raw outputs and run logs were not available with the inspected manuscript; no independent regeneration of this campaign is claimed. The conclusions are conditional on the stated configurations and require the original materials for reproducibility verification.

6 | Discussion: Limitations and Counterarguments

6.1 | Limitations

The macro-state construction, archive assignment and empirical transition estimate limit the model to a small reported regime space. Rare transitions can be dominated by a prior, while continuous or changing regimes would require a separately specified classification policy. The campaign provides no measured observation-noise, false-alarm or simulator–plant mismatch study. Missing kinetic constants, initial states, decision bounds, random-stream rules and event counters also limit independent reproduction. Predictive understanding of community dynamics is a broader challenge [26], not an outcome established by this surrogate.

6.2 | Alternative Approaches

Robust design and model-predictive control are meaningful alternatives when the decision problem permits them [9], [32]. Neither is excluded by nutrient switching, and neither was benchmarked here. Evolutionary search may be useful for multimodal simulator objectives, but the present fixed-species continuous design does not establish a combinatorial community-assembly advantage. The restricted evidence supports comparison of the stated configurations rather than a general dominance claim.

7 | Conclusion

This study reports a combined evolutionary search on two changing simulator configurations. On moving peaks, the plain GA has an incumbent-drop metric of 0.103 ± 0.031 compared with AMPERE's 0.293 ± 0.124. On the two-species chemostat surrogate, mean current-best endpoint titer is 4.08 ± 0.17 g/L for AMPERE and 1.55 ± 0.41 g/L for the plain GA; random immigrants and triggered hypermutation have slightly larger displayed means. The ablation summaries suggest that the mutation response deserves further controlled testing, while they do not establish separate benefits of regime memory or anticipation.

Further work should verify the original event and evaluation accounting, provide the complete simulator and paired outputs, test integration accuracy and vary recurrence, contrast and observation noise independently. A finite probe is insufficient as an optimality certificate. Validation on richer community models and process measurements is required before claims about biological adaptation, sustained reactor production or operational deployment can be made.

Data Availability

Original seed-level outputs, event records, simulator parameters and figure-generation data are required to reproduce the reported summaries. Their availability and release location require author confirmation.

Code availability: The source manuscript identifies exp_05.py and exp_05_results.json, with exp_05_f1.png and exp_05_f2.png. These materials were not available with the inspected manuscript; a complete versioned package and access location should be provided before a reproducibility claim is made.

Funding

This research received no specific grant from any funding agency in the public, commercial, or not-for-profit sectors.

Conflicts of Interest

The authors declare no conflicts of interest.

Author Contributions

Ghannadpour, S. F.: Problem formulation, methodology, and writing, original draft. Ferreira, D. A.: software, validation, and writing, review and editing.

Generative-AI Disclosure

The source manuscript reports assistance from a generative AI language model in text and code preparation. The authors should confirm the final disclosure and remain responsible for the accuracy, provenance, and approval of the revised content.

References

[1] Monod, J. (2012). The growth of bacterial cultures. *Selected papers in molecular biology by Jacques Monod*, *139*(606). https://doi.org/10.1146/annurev.mi.03.100149.002103.

[2] Novick, A., & Szilard, L. (1950). Description of the chemostat. *Science*, *112*(2920), 715-716. https://doi.org/10.1126/science.112.2920.715.

[3] Herbert, D., Elsworth, R., & Telling, R. C. (1956). The continuous culture of bacteria; a theoretical and experimental study. *Microbiology*, *14*(3), 601-622. https://doi.org/10.1099/00221287-14-3-601.

[4] Lenski, R. E., & Travisano, M. (1994). Dynamics of adaptation and diversification: a 10,000-generation experiment with bacterial populations. *Proceedings of the national academy of sciences*, *91*(15), 6808-6814. https://doi.org/10.1073/pnas.91.15.6808.

[5] Elena, S. F., & Lenski, R. E. (2003). Evolution experiments with microorganisms: the dynamics and genetic bases of adaptation. *Nature Reviews. Genetics*, *4*(6), 457. https://doi.org/10.1038/nrg1088.

[6] Kussell, E., & Leibler, S. (2005). Phenotypic diversity, population growth, and information in fluctuating environments. *Science*, *309*(5743), 2075-2078. https://doi.org/10.1126/science.1114383.

[7] Thattai, M., & Van Oudenaarden, A. (2004). Stochastic gene expression in fluctuating environments. *Genetics*, *167*(1), 523-530. https://doi.org/10.1534/genetics.167.1.523.

[8] Bailey, J. E., & Ollis, D. F. (2018). *Biochemical engineering fundamentals*. McGraw-Hill. https://modps71.lib.kmutt.ac.th/xmlui/handle/123456789/324

[9] Bastin, G., & Dochain, D. (1990). *On-line estimation and adaptive control of bioreactors* (Vol. 1, p. 379). Amsterdam: Elsevier. https://api.pageplace.de/preview/DT0400.9781483290980_A23885165/preview-9781483290980_A23885165.pdf

[10] Brenner, K., You, L., & Arnold, F. H. (2008). Engineering microbial consortia: A new frontier in synthetic biology. *Trends in biotechnology*, *26*(9), 483-489. https://doi.org/10.1016/j.tibtech.2008.05.004.

[11] Faust, K., & Raes, J. (2012). Microbial interactions: From networks to models. *Nature reviews microbiology*, *10*(8), 538-550. https://doi.org/10.1038/nrmicro2832.

[12] Goldberg, D. E., & Smith, R. E. (1987). Nonstationary function optimization using genetic algorithm with dominance and diploidy. In *Proceedings of the second international conference on genetic algorithms on genetic algorithms and their application* (pp. 59-68). L. Erlbaum Associates Inc. https://doi.org/10.5555/42512.42521

[13] Branke, J. (2002). *Evolutionary optimization in dynamic environments* (Vol. 1). Dordrecht: Kluwer academic publishers. https://doi.org/10.1007/978-1-4615-0911-0.

[14] Grefenstette, J. J. (1992). Genetic algorithms for changing environments. *Parallel problem solving from nature*, *2*, 137-144. https://cir.nii.ac.jp/crid/1573950398811171968

[15] Cobb, H. G., & Grefenstette, J. J. (1993). Genetic algorithms for tracking changing environments. In *Proceedings of the 5th international conference on genetic algorithms* (pp. 523-530). Morgan Kaufmann Publishers Inc. https://doi.org/10.21236/ADA294075.

[16] Blackwell, T., & Branke, J. (2006). Multiswarms, exclusion, and anti-convergence in dynamic environments. *IEEE transactions on evolutionary computation*, *10*(4), 459-472. https://doi.org/10.1109/tevc.2005.857074.

[17] Nguyen, T. T., Yang, S., & Branke, J. (2012). Evolutionary dynamic optimization: A survey of the state of the art. *Swarm and evolutionary computation*, *6*, 1-24. https://doi.org/10.1016/j.swevo.2012.05.001.

[18] Yang, S., Ong, Y. S., & Jin, Y. (2007). *Evolutionary computation in dynamic and uncertain environments*. Springer Science & Business Media. https://doi.org/10.1007/978-3-540-49774-5.

[19] Morrison, R. W., & De Jong, K. A. (1999). A test problem generator for non-stationary environments. In *Proceedings of the 1999 congress on evolutionary computation-CEC99 (CAT. No. 99th8406)* (Vol. 3, pp. 2047-2053). IEEE. https://doi.org/10.1109/cec.1999.785526.

[20] Li, C., Yang, S., Nguyen, T. T., Yu, E. L., Yao, X., Jin, Y., ... & Suganthan, P. N. (2008). *Benchmark Generator for CEC 2009 competition on dynamic optimization*. https://bura.brunel.ac.uk/bitstream/2438/5897/2/Fulltext.pdf

[21] Yang, S. (2005). Memory-based immigrants for genetic algorithms in dynamic environments. In *proceedings of the 7th annual conference on Genetic and evolutionary computation* (pp. 1115-1122). Association for Computing Machinery. https://doi.org/10.1145/1068009.1068196.

[22] Mavrovouniotis, M., Li, C., & Yang, S. (2017). A survey of swarm intelligence for dynamic optimization: Algorithms and applications. *Swarm and evolutionary computation*, *33*, 1-17. https://doi.org/10.1016/j.swevo.2016.12.005.

[23] Andrews, J. F. (1968). A mathematical model for the continuous culture of microorganisms utilizing inhibitory substrates. *Biotechnology and bioengineering*, *10*(6), 707-723. https://doi.org/10.1002/bit.260100602.

[24] Lotka, A. J. (1926). Elements of physical biology. *Science progress in the twentieth century (1919-1933)*, *21*(82), 341-343. https://archive.org/details/elementsofphysic017171mbp

[25] Gause, G. F. (1934). *The struggle for existence*. Baltimore: Williams and Wilkins. https://www.biodiversitylibrary.org/item/23409#page/8/mode/1up

[26] Widder, S., Allen, R. J., Pfeiffer, T., Curtis, T. P., Wiuf, C., Sloan, W. T., ... & Soyer, O. S. (2016). Challenges in microbial ecology: Building predictive understanding of community function and dynamics. *The ISME journal*, *10*(11), 2557-2568. https://doi.org/10.1038/ismej.2016.45.

[27] May, R. M. (1972). Will a large complex system be stable?. *Nature*, *238*(5364), 413-414. https://doi.org/10.1038/238413a0.

[28] Coyte, K. Z., Schluter, J., & Foster, K. R. (2015). The ecology of the microbiome: Networks, competition, and stability. *Science*, *350*(6261), 663-666. https://doi.org/10.1126/science.aad2602.

[29] Zomorrodi, A. R., & Segrè, D. (2016). Synthetic ecology of microbes: Mathematical models and applications. *Journal of molecular biology*, *428*(5), 837-861. https://doi.org/10.1016/j.jmb.2015.10.019.

[30] Mahadevan, R., Edwards, J. S., & Doyle, F. J. (2002). Dynamic flux balance analysis of diauxic growth in Escherichia coli. *Biophysical journal*, *83*(3), 1331-1340. https://doi.org/10.1016/S0006-3495(02)73903-9.

[31] Balaban, N. Q., Merrin, J., Chait, R., Kowalik, L., & Leibler, S. (2004). Bacterial persistence as a phenotypic switch. *Science*, *305*(5690), 1622-1625. https://doi.org/10.1126/science.1099390.

[32] Beyer, H. G., & Sendhoff, B. (2007). Robust optimization–a comprehensive survey. *Computer methods in applied mechanics and engineering*, *196*(33-34), 3190-3218. https://doi.org/10.1016/j.cma.2007.03.003.

[33] Ross, S. M. (2014). *Introduction to probability models*. Elsevier. https://books.google.com/books/about/Introduction_to_Probability_Models.html?id=A3YpAgAAQBAJ

[34] Page, E. S. (1954). Continuous inspection schemes. *Biometrika*, *41*(1/2), 100-115. https://doi.org/10.1093/biomet/41.1-2.100.

[35] Wolpert, D. H., & Macready, W. G. (1997). No free lunch theorems for optimization. *IEEE transactions on evolutionary computation*, *1*(1), 67-82. https://doi.org/10.1109/4235.585893.

[36] Branke, J. (1999, July). Memory enhanced evolutionary algorithms for changing optimization problems. In *proceedings of the 1999 congress on evolutionary computation-CEC99 (Cat. No. 99TH8406)* (Vol. 3, pp. 1875-1882). IEEE. https://doi.org/10.1109/CEC.1999.785502.

[37] Harris, C. R., Millman, K. J., Van Der Walt, S. J., Gommers, R., Virtanen, P., Cournapeau, D., ... & Oliphant, T. E. (2020). Array programming with NumPy. *Nature*, *585*(7825), 357-362. https://doi.org/10.1038/s41586-020-2649-2.

[38] Virtanen, P., Gommers, R., Oliphant, T. E., Haberland, M., Reddy, T., Cournapeau, D., ... & Van Mulbregt, P. (2020). SciPy 1.0: Fundamental algorithms for scientific computing in Python. *Nature methods*, *17*(3), 261-272. https://doi.org/10.1038/s41592-019-0686-2.

[39] Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in science & engineering*, *9*(3), 90-95. https://doi.org/10.1109/mcse.2007.55.

[40] Wilcoxon, F. (1945). Individual comparisons by ranking methods. *Biometrics bulletin*, *1*(6), 80-83. https://doi.org/10.2307/3001968.

[41] S Holm, S. (1979). A simple sequentially rejective multiple test procedure. *Scandinavian journal of statistics*, 65-70. *Scandinavian journal of statistics* 6(2), 65–70. http://www.jstor.org/stable/4615733

[42] Friedman, M. (1937). The use of ranks to avoid the assumption of normality implicit in the analysis of variance. *Journal of the american statistical association*, *32*(200), 675-701. https://doi.org/10.1080/01621459.1937.10503522.

[43] Demšar, J. (2006). Statistical comparisons of classifiers over multiple data sets. *Journal of machine learning research*, *7*, 1-30. https://doi.org/10.5555/1248547.1248548

[44] Lamtar Gholipoor, M., Alimoradi, M., & Fakheri, S. (2024). A novel metaheuristic approach inspired by trees social relationships and models for fermentation medium. *Metaheuristic Algorithms with Applications, 1 (1), 1-11*. https://doi.org/10.22105/maa.v1i1.17.