Computing, Telecommunication and Control, 2026, Vol. 19, No. 2, Pp. 95–108.
Информатика, телекоммуникации и управление. 2026. Том 19, № 2. С. 95–108.

Research article
DOI: https://doi.org/10.18721/JCSTCS.19209
UDC 004.891.3:519.853

STRUCTURE-AWARE DECISION SUPPORT SYSTEM FOR MOO-DRIVEN CONTROL

A. Deeb ✉ , *V.N. Khokhlovskiy*

Peter the Great St. Petersburg Polytechnic University,
St. Petersburg, Russian Federation

✉ dib_a@spbstu.ru

Abstract. Cyber-physical control systems must satisfy strict constrains on the magnitude of control actions and the rate of signal change, while simultaneously balancing multiple, often conflicting, control objectives. In multi-objective optimization (MOO) problems, the optimizer naturally produces a set of competing control candidates. However, in closed-loop operation, a single implementable control action must be selected at each sampling instant. In practice, this selection is typically performed through fixed scalarization rules, which can lead to performance degradation under changing operating conditions in the case of non-convex Pareto fronts or poorly scalable objectives. This work aims to improve the quality of control based on MOO through a decision support system that accounts for the structure: the proposed architecture decouples online Pareto-front construction from action selection. A heuristic multi-objective optimization algorithm based on the NSGA-II genetic algorithm generates feasible control sequences under magnitude and rate constraints, while a learned selector, denoted as Learned Pareto Filter (LPF), evaluates these candidates using structural features that account for both individual characteristics of the solutions and their interrelations within the Pareto set. The method is evaluated on two benchmark problems. First, a system with an analytically known Pareto front structure is used. Second, the approach is validated on a nonlinear boiler-turbine load-changing scenario, widely used in constrained nonlinear control studies. The results demonstrate that LPF-based selection produces smoother control actions, improved tracking performance, and greater robustness to regime changes compared to objective-based scalarization methods. The findings confirm that, in Pareto-based control, closed-loop quality depends critically on the selection mechanism rather than solely on the optimizer generating the Pareto set.

Keywords: multi-objective optimization, model predictive control, decision support systems, Pareto selection, graph neural networks, constrained nonlinear control

Acknowledgements: The research was financially supported by the Ministry of Science and Higher Education of the Russian Federation within the framework of the program “World-class Research Center: Advanced Digital Technologies” (Agreement No. 075-15-2022-311 dated 20.04.2022).

Citation: Deeb D.A., Khokhlovskiy V.N. Structure-aware decision support system for MOO-driven control. Computing, Telecommunications and Control, 2026, Vol. 19, No. 2, Pp. 95–108. DOI: 10.18721/JCSTCS.19209

© Deeb D.A., Khokhlovskiy V.N., 2026. Published by Peter the Great St. Petersburg Polytechnic University

Научная статья
DOI: https://doi.org/10.18721/JCSTCS.19209
УДК 004.891.3:519.853

СТРУКТУРНО-ОРИЕНТИРОВАННАЯ СИСТЕМА ПОДДЕРЖКИ ПРИНЯТИЯ РЕШЕНИЙ ДЛЯ УПРАВЛЕНИЯ, ОСНОВАННОГО НА МНОГОКРИТЕРИАЛЬНОЙ ОПТИМИЗАЦИИ

А. Диб ✉ , *В.Н. Хохловский*

Санкт-Петербургский политехнический университет Петра Великого,
Санкт-Петербург, Российская Федерация

✉ dib_a@spbstu.ru

Аннотация. Киберфизические системы управления должны удовлетворять строгим ограничениям по величине управляющих воздействий и скорости изменения сигнала, одновременно обеспечивая баланс между несколькими (часто противоречивыми) целями управления. В задачах многокритериальной оптимизации оптимизатор естественным образом генерирует множество конкурирующих кандидатов на управление. Однако в замкнутом контуре на каждом шаге дискретизации требуется выбрать единственное реализуемое управляющее воздействие. На практике этот выбор обычно выполняется с помощью фиксированных правил скаляризации, что может приводить к ухудшению характеристик при изменении режимов эксплуатации в случае невыпуклой формы фронта Парето или трудно масштабируемых критериях. Данная работа направлена на повышение качества управления, основанного на многокритериальной оптимизации, на базе системы поддержки принятия решений, учитывающей структуру: предложенный метод разделяет этап онлайн-построения фронта Парето и этап выбора управляющего воздействия. Эвристический алгоритм многокритериальной оптимизации, основанный на генетических алгоритмах NSGA-II, генерирует допустимые последовательности управляющего воздействия с учетом ограничений по величине и скорости изменений, тогда как предложенный обучаемый селектор, обозначенный как Learned Pareto Filter (LPF), оценивает этих кандидатов, используя структурные признаки, учитывающие как индивидуальные характеристики решений, так и их взаимосвязи внутри множества Парето. Метод оценивается на двух эталонных задачах. Во-первых, используется система с аналитически известной структурой фронта Парето. Во-вторых, подход проверяется на нелинейной задаче изменения нагрузки котлотурбинной установки, широко используемой в исследованиях управления нелинейными системами с ограничениями. Результаты показывают, что выбор на основе LPF обеспечивает более гладкие управляющие воздействия, улучшенную точность слежения и большую устойчивость к изменениям режима по сравнению с методами, основанными исключительно на скаляризации критериев. Полученные выводы подтверждают, что качество Парето-ориентированного управления замкнутым контуром в критической степени определяется механизмом выбора, а не только оптимизатором, формирующим множество Парето.

Ключевые слова: многокритериальная оптимизация, модельно-прогностическое управление, системы поддержки принятия решений, отбор по Парето, графовые нейронные сети, управление нелинейными системами с ограничениями

Финансирование: Исследование выполнено при финансовой поддержке Министерства науки и высшего образования Российской Федерации в рамках программы «Научный центр мирового уровня: Перспективные цифровые технологии» (договор № 075-15-2022-311 от 20.04.2022).

Для цитирования: Deeb D.A., Khokhlovskiy V.N. Structure-aware decision support system for MOO-driven control // Computing, Telecommunications and Control. 2026. Т. 19, № 2. С. 95–108. DOI: 10.18721/JCSTCS.19209

© Диб А., Хохловский В.Н., 2026. Издатель: Санкт-Петербургский политехнический университет Петра Великого

Introduction

New-generation cyber-physical systems are increasingly required to operate within strict safety envelopes, under hard actuator bounds and rate constraints, while simultaneously satisfying multiple, and often conflicting, performance objectives. This requirement is particularly pronounced in energy and process industries, where strong variable coupling, nonlinear dynamics, and continuously shifting operating regimes render controller design a nontrivial task. Thermal power units and boiler-turbine systems constitute representative examples: they must guarantee safety-critical constraints while delivering the demanded power output despite nonlinear interactions and multivariable dependencies [1–3].

Such characteristics naturally motivate control paradigms that explicitly incorporate constraints and performance trade-offs, rather than reducing the problem to a single aggregated objective. In this context, Model Predictive Control (MPC) has emerged as a principled framework for constraint-aware decision making and has been widely adopted in industrial practice [4, 5]. Nevertheless, in highly nonlinear plants, real-time feasibility often necessitates structural simplifications, including local linearization, gain scheduling, or fuzzy approximations around operating points [6–9]. Alternative approaches based on learning-based controllers and evolutionary optimization offer improved adaptability to nonlinearities and uncertainties, yet they frequently introduce considerable computational overhead, which complicates their deployment in strict real-time environments [10–15]. Consequently, practical control design becomes a balance between model expressiveness (nonlinear representations, multi-objective formulations, hard constraints) and the need for immediate actionable decisions.

A common strategy for handling multiple objectives is scalarization, where several performance criteria are combined into a single cost function via weighted sums or related aggregation rules. Although convenient, scalarization requires the selection of preference weights *a priori* and may become fragile when the appropriate trade-off varies across operating regimes [16]. In online control settings, this leads to what may be termed a meta-decision gap: even if a multi-objective optimizer produces a set of nondominated solutions, the closed-loop implementation requires selecting exactly one control input at each sampling instant. If the selection mechanism does not reflect the current regime, the resulting closed-loop behavior may degrade, even when the computed Pareto set is accurate.

To address this meta-decision gap, this paper proposes a structure-aware Decision Support System (DSS) for multi-objective optimization (MOO)-driven control. The proposed DSS operates within a receding-horizon framework. At each step, it constructs an approximate Pareto front of feasible control sequences that satisfy hard input bounds and slew-rate constraints, and subsequently selects one candidate using a learned, structure-exploiting selection mechanism. Instead of relying on fixed scalarization, Pareto selection is formulated as a context-dependent decision problem over a variable-sized set of alternatives. The selector leverages both per-candidate attributes and relational information among candidates – such as dominance relations, smoothness properties, and similarity measures – to adaptively determine the applied control action.

The main contributions of this work can be summarized as follows:

- A set-based formulation for online MOO-driven control under hard input bounds and rate constraints, where the optimizer returns a feasible Pareto set of control sequences and a dedicated selector determines the applied input.
- A structure-aware selection mechanism that evaluates Pareto candidates using features capturing both individual and relational properties, enabling adaptive and regime-dependent trade-off selection without manual weight tuning.
- An experimental validation based on an online Pareto-front construction procedure (NSGA-II) [17], with comparisons against standard scalarization-based selection rules, demonstrating improved decision consistency under changing operating conditions.

Related works

Constraint handling is central in real systems due to actuator saturation and rate limits. MPC provides a principled framework to embed such constraints directly in control synthesis [4–5]. In boiler-turbine and related thermal systems, constrained predictive control has often relied on model simplifications, including fuzzy multi-model schemes and operating-point strategies [3, 6–7], as well as robust and nonlinear variants built on local models or auxiliary constructions to preserve feasibility and stability [18–19]. Online linearization and piecewise representations remain common tools to balance accuracy and computational tractability across regimes [8–9].

When objectives are nonconvex, discontinuous, or expensive to evaluate, evolutionary computation offers flexible black-box optimization and has been applied to control design and online tuning [13–14]. Learning-based controllers can further capture nonlinearities and uncertainties, improving performance over classical baselines [10–11, 20]. However, real-time deployment is constrained by training, inference, and repeated optimization costs within the loop [12], motivating architectures that separate expensive policy search from fast online execution while preserving constraint awareness and regime adaptability.

MOO naturally arises in constrained control problems involving competing criteria such as tracking, energy usage, and safety margins [16]. Evolutionary multi-objective algorithms are widely used to approximate Pareto fronts in complex spaces [17], yet they do not resolve the core online requirement: selecting a single admissible action from a set of nondominated candidates. In practice, scalarization is the dominant baseline, but it encodes fixed global preferences, can miss nonconvex regions, and is sensitive to scaling and regime changes [16, 21].

Theoretical works on multi-objective MPC analyze how selection rules affect stability and performance. Formal multi-objective MPC formulations link Pareto concepts with receding-horizon control [22], utopia-based selection has been used to reason about stability [23], and hierarchical/lexicographic prioritization has been studied to avoid destabilizing trade-offs in nonlinear MPC [24]. Moreover, it has been shown that closed-loop guarantees may be achieved without approximating the entire Pareto front at each step, provided that an appropriate Pareto-optimal sequence is selected [25]. This highlights the importance of the selection mechanism itself.

From a decision-support and multi-criteria decision-making (MCDM) perspective, choosing among Pareto candidates depends on scaling, geometry, and operating context, and fixed scalarizations are often fragile in practice [21]. Recent learning-based approaches, including multi-objective reinforcement learning (MORL), treat Pareto coverage and online selection as primary objectives, guiding policy search with predictive signals to improve front quality and adaptability [26]. However, such methods typically do not emphasize explicit constraint feasibility and receding-horizon structure required in hard real-time systems.

Taken together, the literature reveals a persistent gap: classical multi-objective MPC provides constraint-aware predictive structure and stability analysis under fixed selection principles [21–25], while learning-based multi-objective control emphasizes adaptive selection but often abstracts away strict constraint handling [26]. A structure-aware decision support system for MOO-driven control lies between these paradigms: it preserves explicit constraint feasibility and short-horizon predictive optimization, yet replaces fragile fixed scalarizations with learned, context-aware selectors that operate directly on the structure of the Pareto candidate set.

Problem formulation

Consider a discrete-time controlled system:

$$x_{k+1} = f\left(x_k, u_k\right), \quad x_k \in \mathbb{R}^n, \; u_k \in \mathbb{R}^m, \tag{1}$$

operating under hard input bounds and rate (slew) limits:

$$u_{\min} \le u_k \le u_{\max}, \quad \Delta u_{\min} \le u_k - u_{k-1} \le \Delta u_{\max}, \tag{2}$$

where inequalities are elementwise. At each time step, the controller must choose an admissible action that balances multiple, potentially conflicting objectives. Let

$$g(x) = \begin{bmatrix} g_1(x) \\ g_2(x) \\ \vdots \\ g_p(x) \end{bmatrix} \in \mathbb{R}^p \tag{3}$$

denote a vector of performance criteria to be minimized. For a horizon H, we define a candidate control sequence:

$$U = (u_0,\ u_1, \ldots,\ u_{H-1}) \in \mathbb{R}^{Hm}, \tag{4}$$

and its predicted terminal state:

$$x_H = \Phi_H(x_k, U); \quad \Phi_H(x_k, U) := f(\cdot, u_{H-1}) \circ \cdots \circ f(\cdot, u_0)(x_k). \tag{5}$$

The induced multi-objective cost of the candidate is then:

$$J(U; x_k) := g(x_H). \tag{6}$$

To reduce actuator chattering and encourage smooth actions we add a quadratic increment penalty

$$S(U; u_{k-1}) := \sum_{i=0}^{H-1} \|u_i - u_{i-1}\|_2^2, \quad u_{-1} := u_{k-1}. \tag{7}$$

The online decision problem at time k can be stated as a constrained MOO over sequences:

$$\min_U \left(J(U; x_k), \lambda_s S(U; u_{k-1}) \right), \tag{8}$$

where U satisfies (2) across the horizon, and $\lambda_s \ge 0$ trades off smoothness.

The meta-decision gap in Pareto-based control

A central issue in Pareto-based control is that the optimizer returns alternatives, but the closed loop requires one actionable control input. In many Multiple-Input and Multiple-Output (MIMO) applications, it is common to collapse the objective vector into a single scalar cost via a fixed linear combination (weighted sum),

$$\tilde{J}(U_k) = \sum_{j=1}^{q} w_j J^{(j)}(U_k; z_k), \quad w_j \ge 0. \tag{9}$$

While convenient, such scalarizations implicitly assume that preferences are known *a priori* and remain valid across operating regimes. Moreover, linear scalarizations can fail to represent important trade-offs and often require extensive retuning when conditions change. In practice, this “merge-the-

objectives" step becomes a bottleneck: even if $\mathcal{P}_k$ is accurate, a poor selection rule can degrade closed-loop performance.

Learning-based preference modeling

The proposed strategy is to learn a selector π_0 offline from preference labels generated by an oracle utility $u^*(\cdot)$:

$$u_k^*(i) = u^*\left(c_k^{(i)}; s_k\right), \tag{10}$$

designed to reflect tolerance-aware preferences. A function approximator is trained to predict $\hat{u}_k(i)$ from $(\mathcal{G}_k, i)$, and the online decision is

$$i^* = \arg\max_{i \in \mathcal{V}_k} \hat{u}_k(i), \quad u_k = u_{k|k}^{\left(i^*\right)}. \tag{11}$$

This decoupled architecture replaces fixed scalarization by adaptive, data-driven selection while preserving the underlying multi-objective optimizer. Rather than collapsing to a single scalar objective. We approximate the Pareto set of nondominated candidates at each step. Specifically, we define Pareto dominance in objective space: for two candidates U^a and U^b,

$$\begin{aligned} U^a \prec U^b \Leftrightarrow J\left(U^a; x_k\right) \le J\left(U^b; x_k\right) \ \text{ elementwise and} \\ \exists j : J_j\left(U^a; x_k\right) < J_j\left(U^b; x_k\right). \end{aligned} \tag{12}$$

The (approximate) Pareto front at time k is then the set of candidates not dominated by any other candidate in the population:

$$\mathcal{P}(x_k) := \left\{U \in \mathcal{U}_H(x_k) : \nexists \tilde{U} \in \mathcal{U}_H(x_k) \text{ s.t. } \tilde{U} \prec U\right\}, \tag{13}$$

where $\mathcal{U}_H(x_k)$ denotes the set of feasible sequences under (2). Given $\mathcal{P}(x_k)$, a selector policy must choose one candidate to apply in receding-horizon fashion:

$$U_k^* = \pi\left(\mathcal{P}(x_k), x_k, u_{k-1}\right), \quad u_k = \left(U_k^*\right)_0. \tag{14}$$

Graph neural networks as Pareto-front selectors

A key difficulty in (14) is that the Pareto front is a set whose size and composition vary over time. Hand-designed scalarizations such as weighted-sum and Chebyshev methods impose a fixed preference structure and can behave poorly when the front is nonconvex, poorly scaled, or when the appropriate trade-off changes with the state.

Graph neural networks (GNNs) offer a natural way to map a variable-size set of candidates to scores while exploiting relational structure. A GNN operates on a graph $G = (V, E)$ with node features $\{h_v\}_{v \in V}$ and edge features $\{e_{uv}\}_{(u, v) \in E}$, and updates node embeddings through message passing:

$$m_v^{(\ell)} = \sum_{u \in \mathcal{N}(v)} \psi^{(\ell)}\left(h_v^{(\ell)}, h_u^{(\ell)}, e_{uv}\right); \tag{15}$$

$$h_v^{(\ell+1)} \& = \phi^{(\ell)}\left(h_v^{(\ell)}, m_v^{(\ell)}\right), \tag{16}$$

for layers $\ell = 0,\ldots, L-1$, where $\mathcal{N}(v)$ denotes neighbors of v. A permutation-invariant readout then produces a graph-level embedding

$$z(G) = \rho\left(\left\{h_v^{(L)}\right\}_{v\in V}\right), \tag{17}$$

and finally, a scalar score $s(G)$ is computed by a small predictor head. This architecture is well-suited to Pareto selection because:

1) the candidate set is naturally represented as nodes,
2) pairwise relations can be encoded as edges,
3) invariance to candidate ordering is required.

LPF filtering under feasibility constraints

We propose a learned selector, denoted as Learned Pareto Filter (LPF), that chooses u_k by scoring Pareto candidates using a GNN. At time k, a multi-objective evolutionary search (NSGA-II) produces an approximate front

$$\hat{\mathcal{P}}(x_k) = \left\{U^{(1)},\ldots, U^{(N_k)}\right\}, \tag{18}$$

with associated per-candidate attributes

$$\left(x_H^{(i)},\ J^{(i)} = g\left(x_H^{(i)}\right),\ S^{(i)} = S\left(U^{(i)}; u_{k-1}\right)\right). \tag{19}$$

We build a graph G_k whose nodes correspond to candidates $U^{(i)}$. Node features encode both candidate-specific and context features:

$$h_i = \varphi\left(x_k,\ u_{k-1},\ x_H^{(i)},\ J^{(i)},\ S^{(i)}\right), \tag{20}$$

and edges encode candidate-to-candidate relations, such as dominance indicators, objective-space distances, or normalized differences in terminal states:

$$e_{ij} = \eta\left(J^{(i)} - J^{(j)},\ x_H^{(i)} - x_H^{(j)}\right). \tag{21}$$

The LPF model produces a scalar score for each node i:

$$s_i = GNN_\theta\left(G_k; i\right), \tag{22}$$

and selects the candidate with minimal score:

$$i^* = \arg\min_{i\in\{1,\ldots,N_k\}} s_i,\ u_k = \left(U^{(i^*)}\right)_0. \tag{23}$$

Training signal

To ground the learned selector with a concrete notion of progress, we train θ to imitate a teacher that is available in simulation. For the benchmark considered in this work, the teacher evaluates a candidate by the one-step improvement toward the true Pareto set in decision space, optionally with a smoothness term:

$$y_k^{(i)} = d_{\mathcal{P}}\left(x_{k+1}^{(i)}\right) + \alpha\left\|u_0^{(i)} - u_{k-1}\right\|_2^2,\ x_{k+1}^{(i)} = f\left(x_k, u_0^{(i)}\right), \tag{24}$$

where $d_{\mathcal{P}}(\cdot)$ is a known distance-to-Pareto-set metric in the benchmark and $\alpha \geq 0$ controls smoothness. The training objective is mean-squared error

$$\min_{\theta} \mathbb{E}\left[\left(GNN_{\theta}\left(G_k; i\right) - y_k^{(i)}\right)^2\right], \tag{25}$$

evaluated over graphs produced by rollouts in which the teacher closes the loop. At test time, only $\hat{\mathcal{P}}\left(x_k\right)$ and the learned model are used.

Experiments

Decision support effect towards a known Pareto front

To evaluate the decision-support capability of the proposed selector independently of plant complexity, we adopt a controlled benchmark with analytically known Pareto structure. The system is a discrete-time integrator:

$$x_{k+1} = x_k + T_s u_k, \quad x_k \in \mathbb{R}^2, \quad u_k \in \mathbb{R}^2, \tag{26}$$

subject to elementwise magnitude and slew-rate constraints

$$u_{\min} \leq u_k \leq u_{\max}, \quad \Delta u_{\min} \leq u_k - u_{k-1} \leq \Delta u_{\max}. \tag{27}$$

Two conflicting quadratic objectives are defined at the terminal predicted state:

$$g_1\left(x\right) = x_1^2 + x_2^2; \tag{28}$$

$$g_2\left(x\right)\& = \left(x_1 - 1\right)^2 + \left(x_2 - 1\right)^2. \tag{29}$$

Minimizing g_1 attracts the state toward the origin, whereas minimizing g_2 attracts it toward (1, 1). The Pareto-optimal set in decision/state space is analytically known as

$$\mathcal{P}_x = \left\{\left(t,t\right) \middle| t \in \left[0,1\right]\right\}, \tag{30}$$

which provides an exact reference for measuring convergence in both state and objective spaces.

At each time step, an approximate Pareto set over horizon-H admissible control sequences $H = (u_0, \ldots, u_{H-1})$ is constructed using NSGA-II under constraints. Each candidate is evaluated by:

$$\begin{gathered} J\left(U; x_k\right) = \begin{bmatrix} g_1\left(\Phi_H\left(x_k, U\right)\right) \\ g_2\left(\Phi_H\left(x_k, U\right)\right) \end{bmatrix}; \\ S\left(U; u_{k-1}\right) = \sum_{i=0}^{H-1} \left\| u_i - u_{i-1} \right\|_2^2. \end{gathered} \tag{31}$$

From this candidate set, four selection strategies are compared:

1. WS: fixed weighted sum scalarization.
2. CHEB: Chebyshev scalarization relative to a utopia point.
3. PROJ: oracle baseline selecting the candidate whose terminal state is closest to $\mathcal{P}_x$.
4. LPF (ours): learned Pareto filter using a graph neural network.

Only the first control action of the selected candidate is applied in receding-horizon fashion.

Convergence is measured in both decision and objective spaces. The decision-space distance to the Pareto set is

$$d_x(x) = \min_{p \in \mathcal{P}_x} \|x - p\|_2, \tag{32}$$

while objective-space proximity is measured by distance to a dense sampling of the true front $\mathcal{P}_f$:

$$d_f(f) = \min_{q \in \mathcal{P}_f} \|f - q\|_2, \quad f = g(x). \tag{33}$$

Closed-loop simulations are run for T steps from the same initial condition x_0, using identical NSGA-II parameters across all strategies. Training data are generated from simulated rollouts. At each step, NSGA-II provides a candidate Pareto set. For each candidate, a teacher target is computed based on one-step progress toward $\mathcal{P}_x$:

$$y_k^{(i)} = d_x\left(x_k + T_s u_0^{(i)}\right) + \alpha \left\| u_0^{(i)} - u_{k-1} \right\|_2^2. \tag{34}$$

A graph representation of each candidate, encoding objective values, smoothness, and contextual state information, is used to train a GNN to regress this target via mean-squared error. During data generation, the oracle PROJ selector is used to ensure coverage of relevant regions of the state space. Table 1 summarizes the convergence statistics obtained from the closed-loop simulations.

Pareto convergence statistics over $T = 40$ steps.

Table 1

Pareto convergence statistics over $T = 40$ steps

Method	Final d_x	Final d_f	Mean d_x	Mean d_f
LPF	0.314	0.642	**0.405**	**0.663**
WS	**0.242**	**0.072**	0.751	1.280
CHEB	1.119	1.455	1.082	1.997
PROJ	0.278	0.767	0.609	1.221

Objective-space trajectories during the first control steps for all selectors, overlaid on the true Pareto front. LPF rapidly aligns its trajectory with the Pareto front, closely matching the oracle-like PROJ behavior, while WS and CHEB deviate due to fixed scalarization biases.

Fig. 1 provides a geometric interpretation of these results. While all strategies start from the same initial condition, their trajectories in objective space differ markedly. LPF rapidly steers the system toward the Pareto front and then evolves along it.

MPC-based control of boiler-turbine unit

To stress-test our strategy, we design the following experiment as a strict ablation of the decision stage in Pareto-front based nonlinear MPC for the boiler-turbine unit. All components of the control framework are kept identical to those reported in [27] except for the rule used to select a single control sequence from the Pareto front. We also decrease the Genetic Algorithm (GA) parameters to keep the Pareto front estimation in its early stages. This helps reduce calculation time but degrades the efficiency of the control strategy, ensuring that any observed performance difference is solely attributable to the Pareto front selection mechanism.

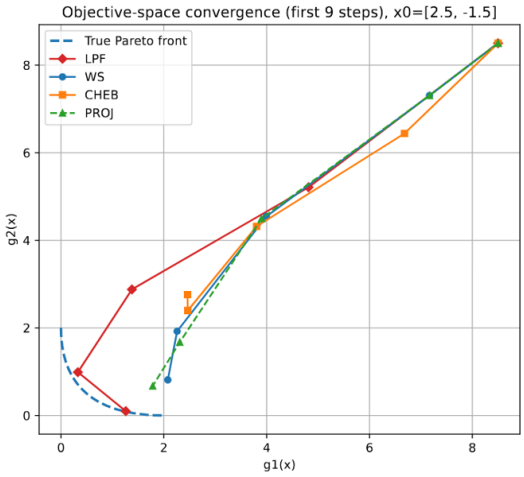


Fig. 1. Objective-space trajectories during the first control steps for all selectors, overlaid on the true Pareto front

The nonlinear model of a boiler-turbine unit is used without approximation or simplification. The state dynamics and the algebraic output relations for $\boldsymbol{y}_1$, $\boldsymbol{y}_2$, $\boldsymbol{y}_3$ are preserved exactly as reported in [27]. The sampling time is $\boldsymbol{T}_s = 0.1$ s. The constraints of magnitude $0 \le u_q \le 1$, $q = 1, 2, 3$ and rate:

$$-0.007 \le \dot{u}_1 \le 0.007, \quad -2 \le \dot{u}_2 \le 2, \quad -0.05 \le \dot{u}_3 \le 0.05,$$

are enforced online through a trimming operator applied to every control update. The system is initialized at the 120 MW operating point:

$$x_{120} = \left[137,\ 120,\ 299.09\right]^T, \quad u_{120} = \left[0{,}563,\ 0.868,\ 0.743\right]^T,$$

and the reference is switched at one quarter of the simulation time to the 140 MW operating point:

$$x_{140} = \left[148.94,\ 140.7,\ 356.3\right]^T.$$

The total simulation time is 400 s. At each policy update, a multi-objective GA searches over a prediction horizon of $H = 10$ steps. The objectives are defined as the discounted absolute tracking errors of the three outputs along the horizon,

$$J_i = \sum_{k=0}^{H-1} \gamma^k \left| y_i(k) - y_i^{des} \right|, \quad i = 1,2,3,$$

with $\gamma = 0.9$. The GA parameters match those of [27]: population size 80, 40 generations, and NSGA-II non-dominated sorting.

Figs. 2 and 3 show the system outputs and control inputs during the load transition. We see that, under the same configurations, the baseline GA fails to control the system near its singular point, and causes a huge mathematical derived from the system equations. As this drift might not be physically accurate, the algorithms are working with the mathematical model that indeed has singular points in the dynamic model. Overcoming this drift is possible by increasing the horizon of the GA, as well as its genetic capability, like population and number of generations. However, our proposed LPF was able to avoid the singularity trap under the same configurations, showing a more stable performance. The main conclusions are:

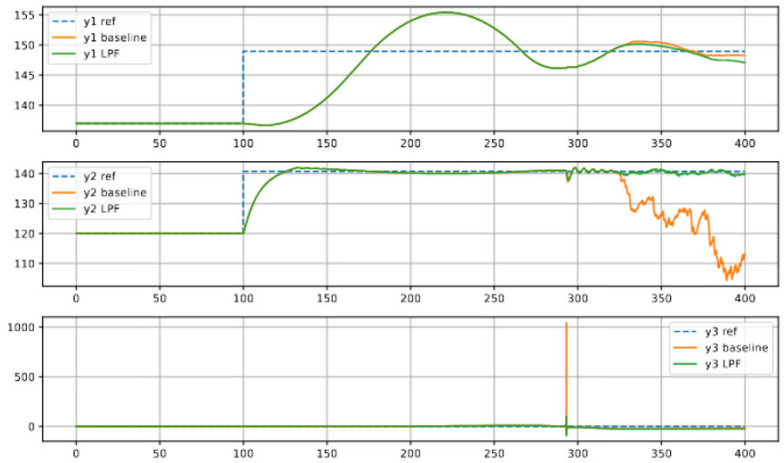


Fig. 2. Comparison between baseline GA selection and LPF-based selection during the MW to MW load transition. The LPF controller achieves smoother inputs and improved tracking performance

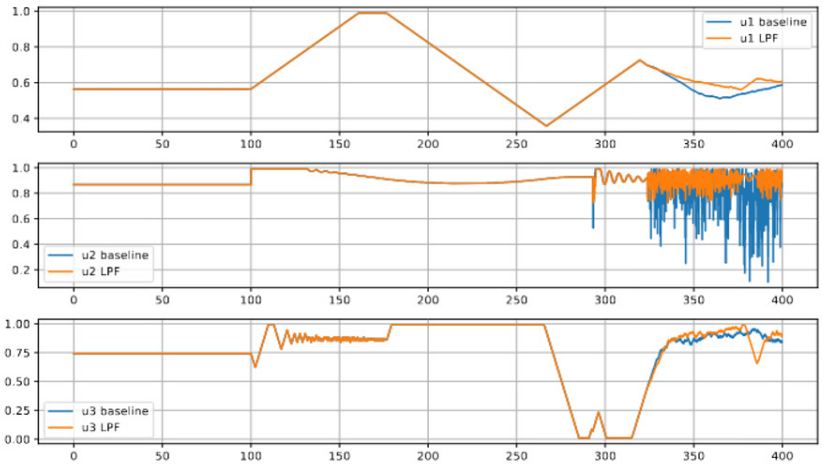


Fig. 3. Comparison between the control signals of the two algorithms

• The LPF-based controller produces significantly smoother control actions, avoiding the oscillatory behavior often selected by the baseline GA policy.

• The tracking of drum pressure, water level, and power output is faster and exhibits reduced overshoot with LPF.

• The LPF policy maintains the system inside the physically valid operating region of the boiler model, preventing the undesirable excursions that appear with purely objective-based selection.

These results demonstrate that the LPF does not alter the optimization process itself, but rather improves the decision stage by selecting physically consistent trajectories from the Pareto front. This confirms that structural awareness of candidate solutions leads to better closed-loop performance than traditional objective-based ranking, even when both methods use the same optimization backbone.

Conclusion

This paper introduced a structure-aware DSS for MOO-driven control under hard magnitude and slew-rate constraints. The core contribution lies in explicitly addressing the meta-decision gap inherent in Pareto-based receding-horizon control: while multi-objective optimizers provide sets of nondominated candidates, practical implementation requires selecting a single admissible control action at each time step. By decoupling Pareto-front generation from action selection, the proposed framework preserves the expressive power of MOO while replacing fragile fixed scalarization rules with a learned, structure-aware selector. The LPF leverages both candidate-specific attributes and relational information within the Pareto set, enabling context-dependent trade-off selection without manual weight tuning. Importantly, the LPF operates as a lightweight module, making it suitable for real-time constrained control architectures.

Experimental results on an analytically tractable benchmark demonstrated that the learned selector consistently guides the system toward the true Pareto set more reliably than classical weighted-sum or norm-based scalarizations. In the nonlinear boiler-turbine load-change scenario, the LPF-based controller achieved smoother control inputs, improved tracking performance, and avoided singular or unstable regions encountered by purely objective-based selection, despite identical optimization backbones. These findings confirm that, in Pareto-based control, the primary performance bottleneck often lies not in Pareto-front approximation but in the selection stage.

Future work will focus on extending the selector to provide theoretical closed-loop guarantees, integrating stability-aware training objectives, and exploring transferability across operating regimes and system classes. Additionally, incorporating uncertainty-awareness and adaptive online refinement of the selector may further enhance robustness in large-scale industrial applications.

REFERENCES

1. IEA. *World Energy Outlook 2022*, Available: Paris https://www.iea.org/reports/world-energy-outlook-2022 (Accessed 06.07.2026).

2. **Röth C., Milde F., Trebbels D., Schmidt J., Doppelbauer M.** A Stator With Offset Segments and a Double Stator Design for the Reduction of Torque Ripple of a Switched Reluctance Motor. *IEEE Transactions on Energy Conversion*. 2022, Vol. 37, No. 2, Pp. 1233–1240.

3. **Wang C., Liu M., Zhao Y., Qiao Y., Chong D., Yan J.** Dynamic modeling and operation optimization for the cold end system of thermal power plants during transient processes. *Energy*, 2018, Vol. 145, Pp. 734–746. DOI: 10.1016/j.energy.2017.12.146

4. **Ogata K.** *Modern control engineering*, 5th ed., New Jersey: Pearson Education, Inc., 2020.

5. **Schwenzer M., Ay M., Bergs T., Abel D.** Review on model predictive control: An engineering perspective. *The International Journal of Advanced Manufacturing Technology*, 2021, Vol. 117, No. 5, Pp. 1327–1349. DOI: 10.1007/s00170-021-07682-3

6. **Wu X., Shen J., Li Y., Lee K.Y.** Stable model predictive control based on TS fuzzy model with application to boiler-turbine coordinated system. *2011 50th IEEE Conference on Decision and Control and European Control Conference*, 2011, Pp. 2981–2987. DOI: 10.1109/CDC.2011.6160553

7. **Wu X., Shen J., Li Y., Lee K.Y.** Hierarchical optimization of boiler-turbine unit using fuzzy stable model predictive control. *Control Engineering Practice*, 2014, Vol. 30, Pp. 112–123. DOI: 10.1016/j.conengprac.2014.03.004

8. **Ławryńczuk M.** Nonlinear predictive control of a boiler-turbine unit: A state-space approach with successive on-line model linearisation and quadratic optimization. *ISA Transactions*, 2017, Vol. 67, Pp. 476–495. DOI: 10.1016/j.isatra.2017.01.016

9. **Yang C., Zhang T., Zhang Z., Sun L.** MLD–MPC for ultra-supercritical circulating fluidized bed boiler unit using subspace identification. *Energies*, 2022, Vol. 15, No. 15, Art. no. 5476. DOI: 10.3390/en15155476

10. **Banzhaf W., Machado P., Zhang M.** *Handbook of evolutionary machine learning*, NY: Springer, 2023.

11. **Lawrence N.P., Damarla S.K., Kim J.W., Tulsyan A., Amjad F., Wang K., Chachuat B., Lee J.M., Huang B., Bhushan Gopaluni R.** Machine learning for industrial sensing and control: A survey and practical perspective. *Control Engineering Practice*, 2024, Vol. 145, Art. no. 105841. DOI: 10.1016/j.conengprac.2024.10-5841

12. **Chuadhry M.A., M.R. G.R., Mathur A.P.** Challenges in machine learning based approaches for real-time anomaly detection in industrial control systems. *Proceedings of the 6th ACM on Cyber-Physical System Security Workshop*, 2020, Pp. 23–29. DOI: 10.1145/3384941.340958

13. **Slowik A., Kwasnicka H.** Evolutionary algorithms and their applications to engineering problems. *Neural Computing and Applications*, 2020, Vol. 32, Pp. 12363–12379. DOI: 10.1007/s00521-020-04832-8

14. **Deeb A., Khokhlovskiy V.N., Shkodyrev V.P.** Model predictive control and genetic algorithms for optimization of continuous stirred tank reactors. In: *Smart electromechanical systems* (eds. I.L. Tarasova, B.A. Kulik), 2024, Vol. 544, Pp. 185–191. DOI: 10.1007/978-3-031-64277-7_14

15. **Biegler L.T.** A perspective on nonlinear model predictive control. *Korean Journal of Chemical Engineering*, 2021, Vol. 38, No. 7, Pp. 1317–1332. DOI: 10.1007/s11814-021-0791-7

16. **Jones D.F., Florentino H.O.** Multi-objective optimization: Methods and applications. In: *The Palgrave handbook of operations research* (eds. S. Salhi, J. Boylan), 2022, Pp. 181–207. DOI: 10.1007/978-3-030-96935-6_6

17. **Deb K., Pratap A., Agarwal S., Meyarivan T.** A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 2002, Vol. 6, No. 2, Pp. 182–197. DOI: 10.1109/4235.996017

18. **Wang L., Cai Y., Ding B.** Robust model predictive control with bi-level optimization for boiler-turbine system. *IEEE Access*, 2021, Vol. 9, Pp. 48244–48253. DOI: 10.1109/ACCESS.2021.3066371

19. **Köhler J., Soloperto R., Müller M.A., Allgöwer F.** A computationally efficient robust model predictive control framework for uncertain nonlinear systems. *IEEE Transactions on Automatic Control*, 2020, Vol. 66, No. 2, Pp. 794–801. DOI: 10.1109/TAC.2020.2982585

20. **Wei Q., Liu Y., Lu J., Ling J., Luan Z., Chen M.** A new integral critic learning for optimal tracking control with applications to boiler-turbine systems. *Optimal Control: Applications and Methods*, 2023, Vol. 44, No. 2, Pp. 830–845. DOI: 10.1002/oca.2792

21. **Peitz S., Dellnitz M.** A survey of recent trends in multiobjective optimal control – surrogate models, feedback control and objective reduction. *Mathematical and Computational Applications*, 2018, Vol. 23, No. 2, Art. no. 30. DOI: 10.3390/mca23020030

22. **Bemporad A., de la Peña D.M.** Multiobjective model predictive control. *Automatica*, 2009, Vol. 45, Pp. 2823–2830. DOI: 10.1016/j.automatica.2009.09.032

23. **Zavala V.M., Flores-Tlacuahuac A.** Stability of multiobjective predictive control: A utopia-tracking approach. *Automatica*, 2012, Vol. 48, No. 10, Pp. 2627–2632. DOI: 10.1016/j.automatica.2012.06.066

24. **He D., Wang L., Sun J.** On stability of multiobjective NMPC with objective prioritization. *Automatica*, 2015, Vol. 57, Pp. 189–198. DOI: 10.1016/j.automatica.2015.04.024

25. **Grüne L., Stieler M.** Multiobjective model predictive control for stabilizing cost criteria. *Discrete and Continuous Dynamical Systems – B*, 2019, Vol. 24, No. 8, Pp. 3905–3928. DOI: 10.3934/dcdsb.2018336

26. **Xu J., Tian Y., Ma P., Rus D., Sueda S., Matusik W.** Prediction-guided multi- objective reinforcement learning for continuous robot control. *Proceedings of the 37th International Conference on Machine Learning*, 2020, pp. 10607–10616.

27. **Deeb A., Khokhlovskiy V., Shkodyrev V.** Hierarchical multi-objective control of nonlinear systems with dynamical input constraints. *Artificial Intelligence and Applications*, 2026, Vol 4., No. 1., Pp. 57–67. DOI: 10.47852/bonviewAIA52024314

INFORMATION ABOUT AUTHORS / СВЕДЕНИЯ ОБ АВТОРАХ

Ali Deeb
Диб Али
E-mail: dib_a@spbstu.ru

Vladimir N. Khokhlovskiy
Хохловский Владимир Николаевич
E-mail: 78v.kh77@gmail.com

Submitted: 28.02.2026; Approved: 14.05.2026; Accepted: 16.06.2026.
Поступила: 28.02.2026; Одобрена: 14.05.2026; Принята: 16.06.2026.