Ensemble Strategies for Objective Subset Selection: A Preference-Aware Consensus Framework

Tomás Marques
Institute for Polymers and Composites
University of Minho
Braga, Portugal
pg50789@alunos.uminho.pt

Sunith Bandaru
School of Engineering Science
University of Skövde
Skövde, Sweden
sunith.bandaru@his.se

António Gaspar-Cunha
Institute for Polymers and Composites
University of Minho
Braga, Portugal
agc@dep.uminho.pt

Abstract— Many-objective optimization in injection molding can involve dozens of strongly dependent quality indicators, which hinders both optimization and result interpretation. This paper proposes a preference-aware ensemble objective-selection framework for conformal cooling channel design. Dependence-based clustering, mRMR, and Laplacian Score are combined through a redundancy-aware consensus mechanism, with optional static decision-maker weights applied at the final ranking stage. The study is positioned as a decision-support-oriented simplification task: the goal is to identify compact, physically interpretable objective subsets for optimization, not to preserve the full 34-objective Pareto structure. The framework is assessed on an injection-mold design case study using a surrogate-assisted NSGA-III workflow. Each methodology yields a four-objective subset (S4) and an associated three-objective subset (S3); the structural loss from S4 to S3 is evaluated in a common S3 space using HV and IGD+. The results show that combining clustering and Laplacian Score provides the most stable reduction, while DM weighting can steer the final subset without dominating the consensus.

Keywords— objective selection, multi-objective optimization, injection molding, conformal cooling channels, mutual information, clustering, artificial neural networks

I. INTRODUCTION

Many real-world engineering optimization problems are naturally formulated with multiple performance indicators, often extending beyond the classical two or three objectives into many-objective settings. However, as the number of objectives increases, optimization and analysis become substantially more difficult: dominance relations lose selection pressure (many solutions become mutually non-dominated), diversity maintenance becomes harder, and both visualization and performance assessment become less informative [1,2].

In this context, selecting a small, informative subset of objectives (e.g., [3,4]) is useful as a decision-support simplification step, helping to identify compact, interpretable formulations without requiring full preservation of the original objective structure. Objective subset selection seeks to remove redundant or weakly informative objectives without discarding the key conflict structure that drives decision making.

Recent work on objective selection/reduction for many-objective optimization can be grouped into several approaches: (i) dominance-preserving methods for non-dominated solutions [3,4]; (ii) unsupervised feature-selection strategies [5]; (iii) comparative analysis of results under reduced objective sets [6]; (iv) data-mining and dimensionality-reduction frameworks [7-10]; and (v) multi-objective reformulations [11]. These methods aim to retain the most informative objectives while improving search efficiency and interpretability.

However, a key limitation of those methods lies in their difficulty of application to real-world problems, which are often characterized by complex relationships among decision variables and objectives.

Motivated by these challenges, this work proposes a preference-aware consensus framework that combines complementary reduction techniques (dependence-based clustering, mRMR, and Laplacian Score) to build compact objective sets for injection-molding optimization. The goal is decision-support-oriented simplification: to identify a small subset of physically interpretable objectives that captures the main trade-offs relevant to the study, rather than to claim full preservation of the original 34-objective Pareto structure.

The main contributions of this work are threefold: i) combining heterogeneous selectors into a redundancy-aware ensemble that outputs compact objective subsets; ii) incorporating static decision-maker preferences as optional weights at the final consensus stage, thereby steering the ranking of candidate objectives without introducing online interactive adaptation; and iii) evaluating the structural loss incurred when moving from a four-objective subset (S4) to its three-objective counterpart (S3) by comparing their projected non-dominated results in a common S3 space.

The paper is organized as follows: Section II introduces the plastics injection molding process; Section III presents the background and the proposed methodology; Section IV presents and discusses the results; and Section V states the conclusions.

II. BACKGROUND

A. *Methods to reduce the number of objectives*

Brockhoff and Zitzler [3,4] introduced two formal optimization problems for objective reduction: the δ-MOSS problem (finding the smallest objective subset that induces an error below a threshold δ) and the k-EMOSS problem (finding a subset of fixed size k with minimum error). Exact and greedy algorithms were proposed to reduce objectives while preserving dominance relations, and validated on knapsack instances and DTLZ benchmarks (DTLZ2, DTLZ5, DTLZ7).

To solve δ-MOSS and k-EMOSS, López et al. [5] proposed an unsupervised feature-selection method using a correlation matrix from the non-dominated set. Objectives are clustered into homogeneous neighborhoods, assuming that greater inter-objective distances indicate stronger conflicts; representative objectives at cluster centers are retained, and the results are compared with [4].

Singh et al. [6] presented the Pareto Corner Search Evolutionary Algorithm (PCSEA), which focuses on identifying Pareto-front corners rather than approximating the full front. Corner solutions help identify relevant objectives and discard redundant ones, with validation on benchmark and engineering problems.

Deb and Saxena [7] proposed a PCA-based reduction framework, later extended to nonlinear variants and ranking-oriented formulations [8–10]. Yuan et al. [11] further formulated objective reduction as a search problem over candidate subsets. These approaches are relevant but may require additional modeling choices or parameter tuning, and they do not necessarily provide a simple consensus view across heterogeneous selectors. Interactive approaches based on online DM feedback have also been proposed [12]. In contrast, the present work is not interactive: DM information, when used, is introduced only as optional static weights in the final consensus stage.

More recently, machine–learning–driven formulations that aim to maintain representativeness under many-objective complexity have been proposed [13]. Additionally, objective selection is not purely a data-driven task: in practice, it must remain consistent with decision-maker priorities and the articulation of preferences. Recent work continues to strengthen preference-based evolutionary multiobjective optimization, including updated performance metrics that more accurately reflect user-preference regions in algorithmic assessments [14].

B. *The injection molding process*

Injection molding is a major manufacturing process for producing plastic parts with high dimensional accuracy, complex geometries, and consistent quality, making it highly suitable for mass production [15]. The molding cycle includes plasticizing (melting and homogenizing polymer pellets), mold closing, filling (high-pressure injection), packing (maintaining pressure to compensate for shrinkage), cooling, and mold opening/ejection [15]. Among these stages, cooling is typically the longest and can dominate the total cycle time, making it essential to improve cooling efficiency for productivity gains while preserving part quality. Injection molding optimization is complex because performance depends on many coupled process variables (e.g., injection/packing conditions, melt and mold temperatures, cooling time) and mold design parameters (e.g., gate and runner geometry, cooling channel layout). The use of conformal cooling channels (CCCs), often enabled by additive manufacturing, enhances thermal control by better matching channel paths to part geometry, improving temperature uniformity and reducing cycle time [16]. However, CCC design also introduces trade-offs in terms of cost, manufacturability, and mold strength.

Because objectives can conflict, injection molding is inherently a multi-objective optimization problem. Common objectives include minimizing cycle time, warpage, and shrinkage while improving thermal uniformity and dimensional stability. A key step is objective selection, focusing on a small set of representative metrics that capture the main productivity, quality trade-offs. Simulation-based optimization combined with surrogate models (e.g., neural networks or response surfaces) enables efficient exploration of these trade-offs with reduced computational cost [16].

C. *Objectives to be addressed*

This study evaluates AI-based multi-objective optimization on an injection mold used at the university to produce a polypropylene cup, providing a realistic benchmark with relevant thermal and quality challenges. The analysis focuses on how different gate locations and two-ring conformal cooling channel (CCC) designs influence part behavior, while all processing conditions are held fixed to ensure that performance changes are driven solely by the cooling design [17]. This implies the use of seven decision variables, as illustrated in Figure 1: the gate location, channel diameters, the distance from the bottom to the channels, and the distance from the channels to the cup. Table I presents the decision variables and their range of variation.

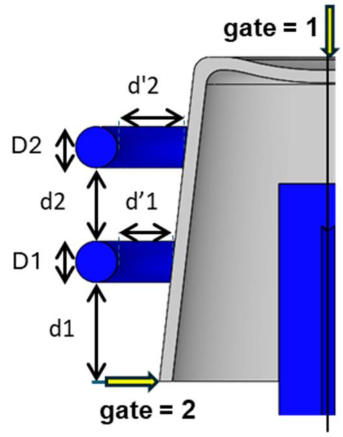

Fig. 1. Two-ring CCC system to be optimized: design variables (d1, d2, D1, D2, d’1, d’2).

To address the high dimensionality and nonlinearity of the problem, the objectives are organized into complementary categories that reflect the main productivity–quality trade-offs. Temperature-based objectives quantify thermal gradients at the end of cooling, including through-thickness and surface-to-core differences. To better capture critical hot spots, the formulation emphasizes the 20% of regions with the highest temperature variation, targeting the areas most prone to residual stresses and deformation. Defect-based objectives assess part integrity using multiple indicators of warpage (maximum and minimum displacements in X, Y, Z, and total deformation) and shrinkage (directional and volumetric changes), where separating maxima and minima helps detect asymmetries caused by flow direction. Time-based objectives address productivity by evaluating cycle time, including Moldex3D-predicted cooling time. Density-based objectives use the density's standard deviation to promote uniform packing and dimensional stability. A thermal efficiency objective (*coolef*) compares the heat removed by the fixed and moving mold halves to quantify the cooling balance. Finally, Von Mises stress objectives (maximum, mean, and standard deviation) evaluate mechanical performance and stress uniformity. Objective selection is supported by dependence-based clustering, mRMR, and Laplacian Score, combined through an ensemble consensus procedure to obtain a compact, interpretable subset.

TABLE I. DECISION VARIABLES AND RANGE OF VARIATION.

Decision variable	Range of variation
d'$_1$ (mm)	[8, 20]
d'$_2$ (mm)	[8, 20]
d1 (mm)	[8, 25]
d2 (mm)	[8, 20]
D1 (mm)	[6, 15]
D2 (mm)	[6, 15]
Gate location	[1, 2]

III. METHODOLOGY PROPOSED

A. Harmony, Conflict, and Correlation

Multi-objective optimization in injection molding, like many other engineering problems, must handle a large set of quality indicators (e.g., cooling time, temperature uniformity, warpage, shrinkage, stresses). These objectives are rarely independent. In practice, they can exhibit harmony, conflict, or redundancy, which directly affects the difficulty of optimization and the value of the computed Pareto set [18].

- Harmonious objectives improve together (e.g., reduced cooling time may reduce thermal gradients in some designs).
- Conflicting objectives exhibit trade-offs (e.g., minimizing cycle time may increase residual stresses).
- Redundant objectives provide nearly the same information, often due to strong dependence or being derived from similar physical mechanisms (e.g., multiple warpage-related measures).

A common way to diagnose redundancy is to quantify inter-objective dependence. In this paper, Pearson correlation is used in the redundancy indicator and blocking rule, whereas Mutual Information (MI) is used to build nonlinear dependence matrices for candidate generation. Spearman correlation is mentioned only for context and is not used in the reported results.

Given two objectives $f_i(\mathbf{x})$and $f_j(\mathbf{x})$, a simple and widely used linear dependence indicator is the Pearson correlation coefficient [19]:

$$\rho(f_i,f_j) = \frac{\mathrm{Cov}\,(f_i,f_j)}{\sigma_{f_i}\sigma_{f_j}} = \frac{\sum_{n=1}^{N}(f_i^{(n)}-\bar{f}_i)(f_j^{(n)}-\bar{f}_j)}{\sqrt{\sum_{n=1}^{N}(f_i^{(n)}-\bar{f}_i)^2}\sqrt{\sum_{n=1}^{N}(f_j^{(n)}-\bar{f}_j)^2}} \quad (1)$$

where N is the number of simulated samples, $\bar{f}_i$ and $\bar{f}_j$ are sample means, and $\sigma_{f_i},\sigma_{f_j}$ are sample standard deviations. Values $\rho \approx 1$or $\rho \approx -1$ indicate a strong positive or negative linear association, while $\rho \approx 0$ suggests weak linear dependence. Although efficient and interpretable, Pearson correlation may underestimate dependence when the relationship between objectives is strongly nonlinear. Rank-based correlation can also be useful when monotonic but nonlinear trends are suspected; however, it is not part of the reported pipeline and is therefore not used in the subsequent selection and evaluation steps.

$$\rho_s(f_i,f_j) = \rho(\mathrm{rank}(f_i),\mathrm{rank}(f_j)) \quad (2)$$

which is robust to nonlinear monotonic trends but may miss more general nonlinear couplings.

To capture broader nonlinear dependence, Mutual Information (MI) is estimated using the Kraskov k-nearest-neighbor estimator [19] with a fixed implementation setting across all reported experiments:

$$I(f_i;f_j) = \sum_a \sum_b p(a,b)\log\left(\frac{p(a,b)}{p(a)p(b)}\right) \quad (3)$$

where $I(f_i;f_j) = 0$ indicates independence (in the ideal limit), and larger values indicate stronger shared information.

When many objectives are dependent, optimizing all of them simultaneously can (i) slow convergence, (ii) inflate computational cost, and (iii) produce Pareto sets dominated by near-duplicate trade-offs. Therefore, objectives should be grouped into highly dependent indicators, from which only a few representative objectives (or latent components) are retained. Conceptually, this corresponds to estimating a dependence structure (correlation or MI matrix), then either:

- reducing dimensionality (latent variables), or
- selecting representative objectives from dependence clusters.

Figure 2 shows the objective dependence map, represented as heatmaps of pairwise Pearson correlation and MI, so that the linear screening view and the nonlinear dependence view can be interpreted side by side.

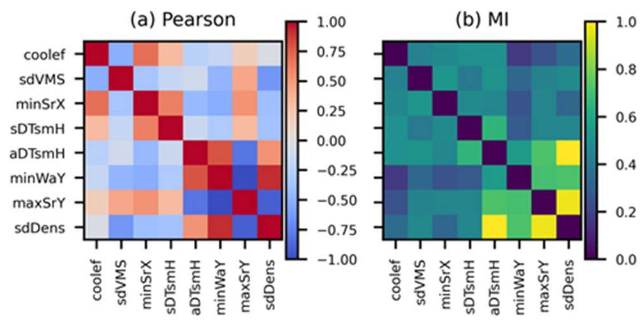

Fig. 2. Heatmap of pairwise dependence: (a) Pearson; and (b) Mutual Information.

B. Selection of Objectives

Injection molding simulations provide multiple quality indicators (e.g., cooling efficiency, temperature uniformity, stresses, shrinkage, density, warpage), leading to 34 optimization objectives. Optimizing all objectives simultaneously is inefficient because many indicators are redundant or highly correlated, thereby increasing computational cost and hindering convergence. Therefore, objective reduction is applied prior to surrogate modeling and many-objective optimization.

We adopt three complementary strategies:

- Clustering-based representative objective selection (dependence-driven clustering);
- Minimum Redundancy–Maximum Relevance (mRMR) (iterative subset selection);
- Laplacian Score (Graph-based Unsupervised Feature Selection)

1) Clustering-based objective selection

Latent components are often difficult to interpret and communicate in industrial settings. For this reason, an alternative strategy selects a reduced set of original objectives while preserving their physical meaning. The procedure begins by computing a dependence matrix between objectives using nonlinear dependence measures such as Mutual Information (MI), HSIC, or Variation of Information [20,21]. Objectives are then clustered, producing sets of highly dependent indicators. From each cluster, one representative objective is selected (e.g., medoid-like choice), yielding a compact and interpretable subset:

$$\mathcal{S}_{\text{cluster}} = \{f_{r_1}, \dots, f_{r_G}\} \tag{4}$$

This approach reduces redundancy while retaining physical interpretability and ensuring that distinct quality aspects remain represented.

2) Minimum Redundancy–Maximum Relevance (mRMR) selection

As a complementary strategy, mRMR is applied in an unsupervised form to build compact subsets that balance informativeness and diversity. Objective relevance can be estimated using variance or entropy, while redundancy is quantified through pairwise dependence. Objectives are selected iteratively to maximize information gain while minimizing overlap with those already chosen [22]:

$$f^\star = \arg\max_{f_i \notin \mathcal{S}} \left(\text{Rel}(f_i) - \beta\,\text{Red}(f_i; \mathcal{S})\right) \tag{5}$$

3) Laplacian Score (Graph-based Unsupervised Feature Selection)

An unsupervised graph-based criterion that ranks objectives by how well they preserve local neighborhood structure in the objective space. A *k*-NN graph is built from the samples, and objectives with smaller Laplacian Scores are preferred; the final subset is obtained by selecting the *K* lowest-score objectives.

$$\Lambda\Sigma(f) = \frac{(\hat{f})^{\mathsf{T}}(D-W)\hat{f}}{(\hat{f})^{\mathsf{T}}D\hat{f}} \tag{6}$$

Together, dependence-based clustering, mRMR, and Laplacian Score provide a robust toolkit for selecting informative, non-redundant, and interpretable objective subsets for many-objective optimization in injection molding.

C. Strategy for Selecting a Reduced Number of Objectives

The clustering-based and mRMR methods described above are applied across multiple parameter settings (e.g., dependence measures, clustering thresholds, subset sizes), producing several candidate subsets:

$$\{\mathcal{S}^{(1)}, \dots, \mathcal{S}^{(T)}\} \tag{7}$$

To consolidate these alternatives into a final reduced set for many-objective optimization, a consensus-based selection strategy is adopted.

1) Consensus indicator

Each objective is assigned a consensus score based on its frequency of occurrence across candidate subsets:

$$C_i = \frac{1}{T}\sum_{t=1}^{T} \mathbb{I}\left(f_i \in \mathcal{S}^{(t)}\right) \tag{8}$$

Optional weights can be provided by the decision maker to favor objectives that are known a priori to be more relevant for the design study. In this work, such weights act only at the final consensus stage; they do not modify the upstream candidate-generation methods and do not represent an interactive adaptation during optimization.

2) Redundancy indicator (correlation-based)

To prevent near-duplicates in the final set, redundancy is computed iteratively using Pearson correlation (after normalization). When growing a selected set $\mathcal{S}$, the redundancy of a candidate objective f_i is measured as its average dependence on already selected objectives:

$$R_i(\mathcal{S}) = \frac{1}{|\mathcal{S}|}\sum_{f_j \in \mathcal{S}} |\rho_{ij}| \tag{9}$$

3) Anchor objective and sequential selection

The procedure begins with an “anchor” objective derived from the global dependency structure. Specifically, the anchor is chosen as the objective with the lowest overall average correlation, ensuring a stable and globally non-redundant starting point. The subset is then expanded sequentially by prioritizing high consensus and low redundancy, using a small augmentation term to break near-ties and ensure deterministic ranking.

To avoid selecting near-duplicates, objectives with a correlation above a predefined threshold ($|\rho| > 0.95$) with the current subset are temporarily blocked. The value 0.95 was chosen as a conservative near-duplicate filter: it excludes almost-collinear objectives while still allowing strongly related but non-redundant objectives to coexist. If all remaining candidates are blocked, the threshold is relaxed in small increments (e.g., +0.01) until at least one candidate becomes admissible.

4) Sensitivity grid and final selection

Since the optimal balance between consensus and redundancy depends on the intended application, a sensitivity grid explores the hyperparameters controlling the emphasis on consensus (λ), redundancy (κ), and the tie-breaking term (μ). For each configuration, the resulting subset and summary indicators (e.g., mean consensus and mean intra-subset correlation) are reported, enabling users to select a final reduced objective set based on interpretability and diversity requirements.

D. Results Assessment

The evaluation has two purposes. First, it assesses whether the ensemble can compress the original 34 indicators into an interpretable four-objective subset (S4), i.e., a first-stage elimination of 30 objectives. Second, it assesses the additional loss incurred when one more objective is removed, yielding S3 (31 eliminated objectives in total). This is a decision-support-oriented simplification test; it does not claim preservation of the full 34-objective Pareto structure.

- Generate and optimize the four-objective formulation S4 for each methodology;
- Generate the associated S3 formulation by removing the fourth-ranked objective from the ordered S4 subset;
- Perform 11 independent NSGA-III runs per formulation under the same search protocol;
- Compare S3 against the projection of S4 into the common S3 space using HV and IGD+ reported separately; the summed score is used only for ranking.

Smaller relative losses indicate that the removed fourth objective may be at least partially represented by the three retained objectives, leading to a more stable S4-to-S3 reduction.

IV. Results and Discussion

A. Problem to Solve and Cases Studied

To evaluate objective subset selection under ensemble voting, we studied three scenarios: (A) selection within each method family (clustering, mRMR, Laplacian score), (B) pairwise combinations between families, and (C) the complete ensemble including decision-maker (DM) weights. This yields seven methodologies. Table II shows the objectives selected for each methodology and case (3- or 4-objective cases). Column S3 shows the 3 objectives selected, and column S4 shows the fourth objective.

TABLE II. Final objective subsets used per methodology

Methodology	S3	S4
A_clust	coolef, sdVMS, minSrX	sDTsmH
A_lap	maxSrY, minWaY, sdVMS	sdDens
A_mrmr	coolef, aDTsmH, sdVMS	minSrX
B_clust_lap	coolef, sdVMS, minSrX	sDTsmH
B_clust_mrmr	coolef, aDTsmH, sdVMS	minSrX
B_mrmr_lap	coolef, aDTsmH, sdVMS	minWaY
C_all	coolef, aDTsmH, sdVMS	minSrX

B. Multi-Output Surrogate Model

A single global multi-output surrogate model was trained to cover both gate configurations, using the gate as an input along with the remaining design variables (7 inputs in total). The model predicts the union of objectives appearing across the studied subsets (8 outputs: coolef, sdVMS, minSrX, sDTsmH, aDTsmH, minWaY, maxSrY, sdDens). After hyperparameter optimization, the final network achieved a mean scaled MAE of approximately 1.71%, with the largest per-output error around 3.08% (minSrX). This accuracy is adequate for repeated runs of many-objective optimization and enables consistent comparison of objective subsets using a fixed surrogate.

C. Optimization Protocol and metric evaluation

All optimization experiments were conducted with NSGA-III using population sizes of 50 and 100, 50 generations, and 11 independent runs per objective set. The two population sizes were included to avoid basing the comparison on a single search budget while keeping the computational effort compatible with the surrogate-assisted study. Candidate solutions were evaluated via vectorized surrogate inference (batch evaluation), which reduces runtime without altering the algorithmic behavior. All objectives were treated as minimization objectives; *coolef* is negative in the dataset, and therefore minimization correctly favors “more negative” values. To quantify the impact of reducing the number of objectives from four to three, metrics must be computed in a consistent space. For each methodology, the results obtained by optimizing S4 were projected into the corresponding S3 Space (denoted S4proj) by retaining only the three objectives shared with S3. Hypervolume (HV) and IGD+ were then computed in the common S3 Space for both S3 and S4proj, allowing a direct assessment of robustness to reduction.

D. Selection of objective results

Because absolute metric magnitudes may vary across cases, robustness was assessed using relative differences with respect to the projected four-objective results (S4proj):

$$\Delta HV_{\text{rel}} = \frac{|HV_{S3}-HV_{S4proj}|}{|HV_{S4proj}|+\varepsilon},\ \Delta IGD^{+}_{\text{rel}} = \frac{|IGD^{+}_{S3}-IGD^{+}_{S4proj}|}{|IGD^{+}_{S4proj}|+\varepsilon} \quad (10)$$

Table III reports the mean relative losses in HV and IGD+ separately, together with a combined ranking score defined as $\Delta HV_{\text{rel}} + \Delta IGD^{+}_{\text{rel}}$ (smaller is more robust). The individual indicators remain the primary evidence; the summed score is used only to order methodologies with similar behavior:

TABLE III. Robustness Ranking

Rank	Methodology	ΔHV_{rel}	$\Delta IGD^{+}_{\text{rel}}$	Score
1	B_clust_lap	0.0493	**0.5694**	0.6187
2	A_clust	0.0500	0.5715	0.6215
3	C_all	**0.0376**	0.6058	0.6434
4	B_mrmr_lap	0.0482	0.6219	0.6701
5	B_clust_mrmr	0.0500	0.6208	0.6708
6	A_mrmr	0.0450	0.6359	0.6809
7	A_lap	0.0866	0.6029	0.6896

In general, B_clust_lap shows the smallest change when moving from S4 to S3, indicating that the reduced three-objective formulation retains a trade-off pattern more similar to the four-objective optimization than the remaining strategies. In contrast, A_lap shows small absolute differences but larger proportional changes due to its lower baseline metric scale, making it less robust in relative comparisons.

E. Assessment Results

To visually support the robustness analysis, the non-dominated solutions from all runs were pooled, and the global non-dominated subset was extracted in the corresponding S3 space. Figure 3 summarizes these pooled sets using parallel coordinates (min–max normalized, best at 0), comparing direct S3 optimization against the projection of the four-objective results onto the same three objectives (S4proj). The close alignment between the S3 and S4proj median profiles suggests that the projected four-objective search reproduces a very similar S3 trade-off pattern, with only limited deviations under the most robust strategy.

Before computing HV and IGD+, all objectives in the common *S3* space are linearly normalized to $[0,1]$using the min–max values taken over the union of nondominated points from $S3$and $S4_{\text{proj}}$(all 11 runs). HV is then computed in the normalized space with reference point $(1.1, 1.1, 1.1)$, and IGD+ is computed using as a reference set the pooled nondominated set obtained from the combined $S3$ and $S4_{\text{proj}}$ outcomes for the same methodology (after normalization).

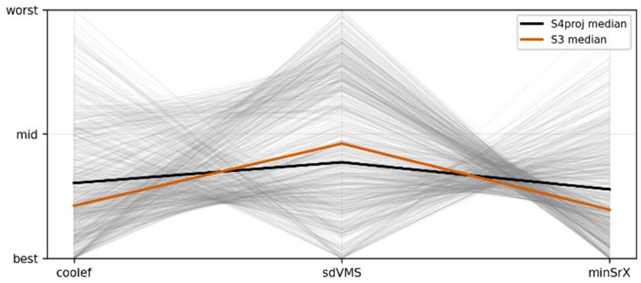

Fig. 3. Combined global non-dominated sets in S3-space represented as parallel coordinates.

V. Conclusions

This work proposed a preference-aware ensemble objective-selection pipeline that combines heterogeneous selectors through a redundancy-aware consensus mechanism. Seven scenarios were analyzed, each yielding a four-objective subset (S4) and its reduced three-objective counterpart (S3). The contribution should be interpreted as a decision-support-oriented simplification strategy, not as a claim of full preservation of the original 34-objective trade-off structure.

Within that scope, the S4-to-S3 analysis indicates that B_clust_lap provides the most stable reduction among the tested methodologies. The results also suggest that the framework is particularly useful in problems with many correlated or partially redundant indicators, where compact subsets can simplify analysis and support more interpretable decision making. Static DM weights can influence the final ranking of objectives, but they do not replace statistical evidence nor provide interactive adaptation during the search.

Future work should strengthen the statistical validation with more independent runs, explore stronger full-space and baseline comparisons, and investigate interactive preference-updating mechanisms.

Acknowledgment

This publication is based upon work from COST Action Randomized Optimization Algorithms Research Network (ROAR-NET), CA22137, supported by COST (European Cooperation in Science and Technology).

References

[1] K. Deb and H. Jain, "An evolutionary many-objective optimization algorithm using reference-point-based nondominated sorting approach, Part I: Solving problems with box constraints," IEEE Trans. Evol. Comput., vol. 18, no. 4, pp. 577–601, 2014, doi: 10.1109/TEVC.2013.2281535.

[2] Y. Tian, L. C. Si, X. Zhang, R. Cheng, C. He, K. C. Tan, and Y. Jin, "Evolutionary large-scale multi-objective optimization: A survey," ACM Comput. Surv., vol. 54, no. 8, pp. 1–34, 2021, doi: 10.1145/3470971.

[3] D. Brockhoff and E. Zitzler, "Are all objectives necessary? On dimensionality reduction in evolutionary multiobjective optimization," in Lecture Notes in Computer Science, 2006, pp. 533–542, doi: 10.1007/11844297_54.

[4] D. Brockhoff and E. Zitzler, "Objective reduction in evolutionary multiobjective optimization: Theory and applications," Evol. Comput., vol. 17, no. 2, pp. 135–166, 2009, doi: 10.1162/evco.2009.17.2.135.

[5] J. A. López, C. C. A. Coello, and D. Chakraborty, "Objective reduction using a feature selection technique," in Proc. 10th Annu. Conf. Genetic and Evolutionary Computation (GECCO '08), 2008, doi: 10.1145/1389095.1389228.

[6] H. K. Singh, A. Isaacs, and T. Ray, "A Pareto corner search evolutionary algorithm and dimensionality reduction in many-objective optimization problems," IEEE Trans. Evol. Comput., vol. 15, no. 4, pp. 539–556, 2011, doi: 10.1109/TEVC.2010.2093579.

[7] K. Deb and D. K. Saxena, "Searching for Pareto-optimal solutions through dimensionality reduction for certain large-dimensional multi-objective optimization problems," in Proc. IEEE Congr. Evol. Comput. (CEC 2006), Vancouver, BC, Canada, 2006, pp. 3353–3360.

[8] D. K. Saxena and K. Deb, "Non-linear dimensionality reduction procedures for certain large-dimensional multi-objective optimization problems: Employing correntropy and a novel maximum variance unfolding," in Evolutionary Multi-Criterion Optimization (EMO 2007), S. Obayashi, K. Deb, C. Poloni, T. Hiroyasu, and T. Murata, Eds. Berlin, Heidelberg: Springer, 2007, vol. 4403, pp. (chapter pagination), doi: 10.1007/978-3-540-70928-2_58.

[9] D. K. Saxena, J. A. Duro, A. Tiwari, K. Deb, and Q. Zhang, "Objective reduction in many-objective optimization: Linear and nonlinear algorithms," IEEE Trans. Evol. Comput., vol. 17, no. 1, pp. 77–99, 2013, doi: 10.1109/TEVC.2012.2185847.

[10] J. A. Duro, D. K. Saxena, K. Deb, and Q. Zhang, "Machine learning based decision support for many-objective optimization problems," Neurocomputing, vol. 146, pp. 30–47, 2014, doi: 10.1016/j.neucom.2014.06.076.

[11] Y. Yuan, Y.-S. Ong, A. Gupta, and H. Xu, "Objective reduction in many-objective optimization: Evolutionary multiobjective approaches and comprehensive analysis," IEEE Trans. Evol. Comput., vol. 22, no. 2, pp. 189–210, 2018, doi: 10.1109/TEVC.2017.2672668.

[12] S. M. Shavarani, M. López-Ibáñez, and R. Allmendinger, "Detecting Hidden and Irrelevant Objectives in Interactive Multiobjective Optimization," IEEE Trans. Evol. Comput., vol. 28, no. 2, pp. 544–557, 2024, doi: 10.1109/TEVC.2023.3269348

[13] A. Gaspar-Cunha, P. Costa, F. Monaco, and A. C. B. Delbem, "Many-objectives optimization: A machine learning approach for reducing the number of objectives," Math. Comput. Appl., vol. 28, no. 1, Art. no. 17, 2023, doi: 10.3390/mca28010017.

[14] D. Yadav, P. Ramu, and K. Deb, "An updated performance metric for preference-based evolutionary multi-objective optimization algorithms," COIN Report no. 2024003, Michigan State Univ., 2024.

[15] T. A. Osswald and J. P. Hernández-Ortiz, Polymer Processing: Modeling and Simulation. Munich, Germany: Hanser, 2006.

[16] S. Feng, A. M. Kamat, and Y. Pei, "Design and fabrication of conformal cooling channels in molds: Review and progress updates," Int. J. Heat Mass Transf., vol. 171, Art. no. 121082, 2021, doi: 10.1016/j.ijheatmasstransfer.2021.121082.

[17] A. Gaspar-Cunha, J. Melo, T. Marques, and A.J. Pontes, "Methodology for Designing Injection Molds: Data Mining and Multi-objective Optimization," in: Applications of Evolutionary Computation. Evo Applications, P. García-Sánchez, E. Hart, and S.L. Thomson, Eds. Cham, Switzerland: Springer, 2025, doi: https://doi.org/10.1007/978-3-031-90062-4_10

[18] T. Chugh, A. Gaspar-Cunha, A. H. Deutz, J. A. Duro, D. C. Oara, and A. Rahat, "Identifying correlations in understanding and solving many-objective optimisation problems," in Many-Criteria Optimization and Decision Analysis, D. Brockhoff, M. Emmerich, B. Naujoks, and R. Purshouse, Eds. Cham, Switzerland: Springer, 2023, doi: 10.1007/978-3-031-25263-1_9.

[19] T. M. Cover and J. A. Thomas, Elements of Information Theory, 2nd ed. Hoboken, NJ, USA: Wiley, 2006.

[20] A. Kraskov, H. Stögbauer, and P. Grassberger, "Estimating mutual information," Phys. Rev. E, vol. 69, no. 6, Art. no. 066138, 2004, doi: 10.1103/PhysRevE.69.066138.

[21] A. Gretton, O. Bousquet, A. Smola, and B. Schölkopf, "Measuring statistical dependence with Hilbert–Schmidt norms," in Algorithmic Learning Theory (ALT 2005), LNCS 3734. Berlin, Germany: Springer, 2005, pp. 63–77, doi: 10.1007/11564089_7.

[22] H. Peng, F. Long, and C. Ding, "Feature selection based on mutual information criteria of max-dependency, max-relevance, and min-redundancy," IEEE Trans. Pattern Anal. Mach. Intell., vol. 27, no. 8, pp. 1226–1238, 2005, doi: 10.1109/TPAMI.2005.159.

.