

ParetoTracker: Understanding Population Dynamics in Multi-objective Evolutionary Algorithms through Visual Analytics

Zherui Zhang, Fan Yang, Ran Cheng, *Senior Member, IEEE*, and Yuxin Ma, *Senior Member, IEEE*

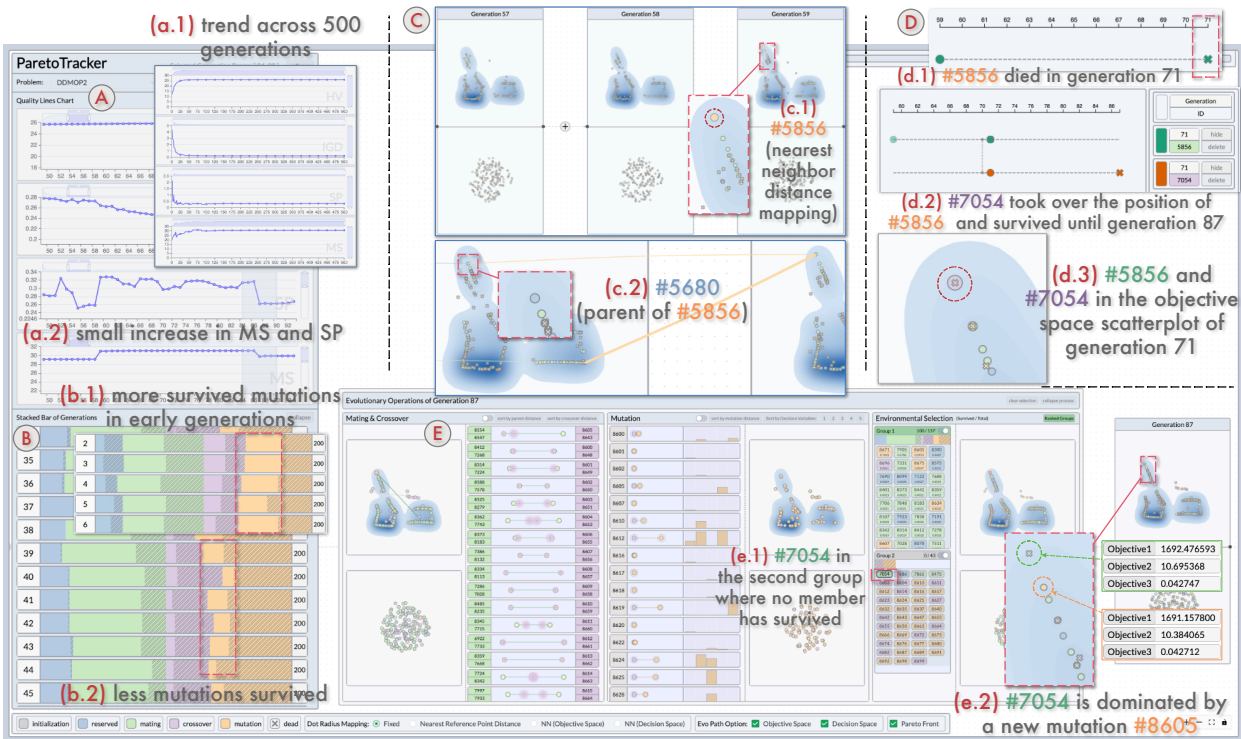


Fig. 1: The interface of ParetoTracker contains five visualization components: (A) quality line charts, (B) stacked bars for generation statistics, (C) decision and objective space scatterplots for selected generations, (D) the lineage view, and (E) the evolutionary operator view.

Abstract—Multi-objective evolutionary algorithms (MOEAs) have emerged as powerful tools for solving complex optimization problems characterized by multiple, often conflicting, objectives. While advancements have been made in computational efficiency as well as diversity and convergence of solutions, a critical challenge persists: the internal evolutionary mechanisms are opaque to human users. Drawing upon the successes of explainable AI in explaining complex algorithms and models, we argue that the need to understand the underlying evolutionary operators and population dynamics within MOEAs aligns well with a visual analytics paradigm. This paper introduces ParetoTracker, a visual analytics framework designed to support the comprehension and inspection of population dynamics in the evolutionary processes of MOEAs. Informed by preliminary literature review and expert interviews, the framework establishes a multi-level analysis scheme, which caters to user engagement and exploration ranging from examining overall trends in performance metrics to conducting fine-grained inspections of evolutionary operations. In contrast to conventional practices that require manual plotting of solutions for each generation, ParetoTracker facilitates the examination of temporal trends and dynamics across consecutive generations in an integrated visual interface. The effectiveness of the framework is demonstrated through case studies and expert interviews focused on widely adopted benchmark optimization problems.

Index Terms—Visual analytics, multi-objective evolutionary algorithms, evolutionary computation

1 INTRODUCTION

- All authors are with the Department of Computer Science and Engineering, Southern University of Science and Technology, China. Email: {zhangzr32021, yangf2020}@mail.sustech.edu.cn, ranchengcn@gmail.com, mayx@sustech.edu.cn. Z. Zhang and F. Yang contributed equally to this work. Y. Ma is the corresponding author.

Manuscript received xx xxx. 201x; accepted xx xxx. 201x. Date of Publication xx xxx. 201x; date of current version xx xxx. 201x. For information on obtaining reprints of this article, please send e-mail to: reprints@ieee.org. Digital Object Identifier: xx.xxx/TVCG.201x.xxxxxxx

Multi-objective optimization problems frequently emerge in various decision-making scenarios such as engineering design [42], environmental planning [3], and artificial intelligence [65]. Unlike single-objective optimization where the goal is to optimize a single criterion, multi-objective optimization considers more than one objectives simultaneously. Since the multiple objectives are often conflicting, improving the performance of one objective might adversely affect others. Consequently, the solution to these problems is not a single optimal point but a set of feasible solutions, termed a *Pareto front*. On this front, each solution is *non-dominated*, indicating that no single solution outperforms others across all objectives consistently.

To address this complexity, the development of multi-objective evolutionary algorithms (MOEAs) has been instrumental. These algorithms leverage evolutionary computation techniques to mediate among the varying objectives [49]. MOEAs have emerged as a pivotal approach for solving multi-objective problems over the past decades, especially noted for their capability to manage complex objective functions and to produce a diverse set of solutions in a single algorithm run. The essential mechanism of MOEAs draws inspiration from natural evolution, where only the fittest solutions in terms of optimization objective functions are preserved [60]. In each generation of the evolutionary process, a population of parent individuals¹ are selected to generate offspring individuals through crossover/mutation operations. Thereafter, an environmental selection step filters these individuals, retaining only those with the highest optimization fitness for the next generation.

Despite the efficacy of MOEAs in resolving complex optimization challenges, their operations often remain opaque, which functions as black-boxes to human users [55]. In the realms of evolutionary computing and visualization, efforts have been made to uncover patterns within optimal solution sets [17]. Nonetheless, facilitating user engagement and understanding of the dynamics in evolutionary processes remains largely unexplored, particularly considering the complex evolution strategies employed by various MOEAs [54, 56]. Conventional methods of reflecting solution characteristics across generations typically involve numerical quality indicators [30], yet a detailed explanation of the evolutionary operations across generations is essential for a comprehensive understanding of how solutions are evolved towards optima. Such insights could empower algorithm developers and practitioners to identify the operational successes or failures of an algorithm as well as gain insights that cannot be gleaned solely through aggregated quality indicators.

The last decade has seen the visualization community successfully apply the visual analytics paradigms to explain complex algorithmic processes, thereby enhancing model transparency [20, 28, 58, 62]. We posit that this paradigm is equally applicable to elucidating the evolutionary dynamics within MOEAs. In this paper, we introduce a visual analytics framework, ParetoTracker, designed to facilitate the understanding and inspection of population dynamics in the evolutionary processes of MOEAs. Drawing from an extensive review of relevant literature and collaboration with domain experts in evolutionary computation, we have derived analytical tasks that inform the visualization and interaction design in a hierarchical manner. By employing the “overview+detail” scheme, ParetoTracker utilizes a multi-level design that exposes the evolutionary progress across generations from three perspectives: the aggregate measures and statistics of solution sets, cross-generational tracking of individuals and their lineages, and in-depth analysis of evolutionary operations between successive generations. The effectiveness of ParetoTracker is demonstrated with case studies and expert interviews on benchmark problems commonly utilized in the MOEA field.

In summary, our contributions include:

- A visual analytics framework, ParetoTracker, for comprehending and inspecting the dynamics of individuals among generations in the evolutionary processes of MOEAs;
- A suite of visual inspection and exploration methods as well as designs to highlight salient evolutionary patterns among generations;
- Case studies and expert interviews on established test problems to demonstrate effectiveness of the framework.

2 RELATED WORK

Our research addresses the challenges involved in examining and investigating the evolutionary processes in MOEAs. This section provides a review of existing literature on the application of visualization techniques in evolutionary multi-objective optimization algorithms as well as explainable AI.

¹To facilitate the explanation in this paper, the terms “individuals” and “solutions” are used interchangeably when discussing evolutionary processes in the context of MOEAs.

2.1 Visualization in Evolutionary Multi-objective Optimization

Visualization is critical in evolutionary multi-objective optimization due to the complexity of decision and objective spaces, typically exceeding two dimensions. High-dimensional visualization techniques are required to represent decision and objective vectors associated with solutions, such as projection techniques [25, 37, 52], parallel coordinates plots [24] (PCPs), and scatterplot matrices [4]. Smedberg and Bandaru [45] propose an interactive visualization environment using linked scatterplot and PCP as the central representation for solutions. Tušar and Filipič’s study [50] on high-dimensional visualization methods highlights their pattern preservation capabilities. They found that projection techniques can support large-scale solution sets and allow simultaneous comparison of multiple sets. However, these techniques may not directly reflect distribution patterns due to potential distortions and two-dimensional result information loss.

The optimization community has also developed specialized techniques for solution set visualization. He and Yen [19] introduce a 2-D radial system for objective vectors. The iSOM method [41, 61] improves interpretability of self-organizing maps. 3D-RadVis [23] shows solution distributions and Pareto front convergence within a single visualization. PaletteViz [47] and PaletteStarViz [46] highlight geometric properties and constraint boundaries closeness in solution sets.

In addition to the static visualization of algorithm outputs, recent efforts have encompassed the visualization of the evolutionary processes to illustrate the dynamics of the algorithms. De Lorenzo et al. [10] and Walter et al. [54, 56] employ multi-dimensional scaling (MDS) to visualize the search trajectory within the decision space for single- and multi-objective optimization problems, respectively, which enables the exploration of complex decision space landscape. In the visual analytics community, VisEvol [8] illustrates the internal crossover and mutation mechanisms in evolutionary algorithms in the context of optimizing machine learning parameters. Huang et al. [22] propose a comparative analysis framework that utilizes timelines and proximity graphs to investigate the similarities between different algorithm runs.

As such, a large body of research has been dedicated to visualizing solution sets and the dynamics of generations. Nonetheless, a limitation among these methods is their static nature, which offers minimal interactive capabilities. Moreover, these approaches often fall short in offering comprehensive tools for detailed analysis of evolutionary processes at varying levels of granularity. To address this gap, our framework enhances analytical capabilities, enabling evaluations from aggregated metrics down to detailed, fine-grained evolutionary operations.

2.2 Visual Analytics in Explainable AI

Over the last decade, a number of surveys [20, 28, 58, 62] have summarized advancements in visual analytics applied to machine learning models. Specifically, the visualization community has proposed various techniques and frameworks for “opening the black box” of machine learning models [2, 15]. Work from Muhlbacher et al. [39] summarizes the strategies for presenting information and integrating user controls in black-box models. Kim et al. [27] introduce the per-iteration visualization environment (PIVE) to provide insights and enable interactive control in the intermediate iterations. In addition to general methodologies, certain studies have concentrated on providing explanations for white-box models, including tree models [33, 34, 51, 66], rule-based representations [38], and linear models [14, 35, 59].

With the evolution of large-scale models such as neural networks, there is a growing necessity to illustrate their training and prediction processes. Usually, the training progress is depicted through statistical charts on a per-iteration basis, which shows key performance metrics like loss or accuracy values [7, 32, 36, 43, 57]. Moreover, there is a body of work dedicated to examining these evolutions through model- or task-specific visualizations, such as DeepTracker [31] for convolutional neural networks, DQNViz [57] for deep reinforcement learning, and GAN Lab [26] for deep generative models.

In summary, the field of explainable AI has witnessed significant research efforts in elucidating the internal mechanisms of complex models. Drawing inspiration from the paradigm of opening black box

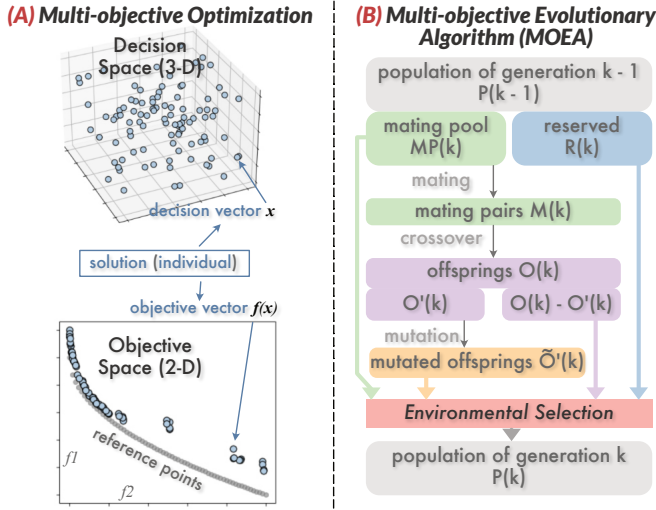


Fig. 2: (A) An example of multi-objective optimization problem. (B) The general evolutionary computing pipeline.

models characterized by their opaque behaviors, our work focuses on the MOEA domain where the evolutionary processes are difficult to comprehend by human analysts. We tend to uncover the complex evolutionary behaviors hidden in their low-level operations, which provides a more detailed perspective to complement the commonly-used aggregated quality indicators [30].

3 BACKGROUND

This section presents a preliminary overview of MOEAs to establish a foundational understanding of the terminology and key components employed in the algorithms.

Multi-objective Optimization. The multi-objective optimization problem can be formulated as minimizing a series of objective functions:

$$\min \mathbf{f}(\mathbf{x}) = (f_1(\mathbf{x}), f_2(\mathbf{x}), \dots, f_m(\mathbf{x})) \quad (1)$$

where \mathbf{x} is an n -dimensional *decision vector* $\mathbf{x} = (x_1, x_2, \dots, x_n)$ in the *decision space* in \mathbb{R}^n , and the function evaluation of \mathbf{x} on m objective functions compose the *objective vector* in the *objective space* in \mathbb{R}^m , denoted as $\mathbf{f}(\mathbf{x})$, Fig. 2 (A). Given that the objectives often show conflicts, e.g., adjusting \mathbf{x} to minimize the value on $f_1(\mathbf{x})$ may increase the outcome on $f_2(\mathbf{x})$, it is unfeasible to find a single solution that simultaneously minimizes all objectives.

For two solutions \mathbf{x}_1 and \mathbf{x}_2 , \mathbf{x}_1 is said to dominate \mathbf{x}_2 (termed $\mathbf{x}_1 \succ \mathbf{x}_2$) if and only if $\mathbf{f}(\mathbf{x}_1)$ is no worse than $\mathbf{f}(\mathbf{x}_2)$ on all objectives and strictly better in at least one objective. A solution \mathbf{x} is deemed *Pareto optimal* if there is no other solution in the objective space that dominates \mathbf{x} . The set of all Pareto optimal solutions forms the *Pareto Set*, while the corresponding objective vectors of these solutions create the *Pareto Front*. Given such definitions, the primary goal of solving multi-objective optimization problems is to obtain the Pareto.

It should be noted that for benchmarking purpose, some artificial test problems offer *reference sets* that contain numerous objective vectors sampled near the Pareto fronts, Fig. 2 (A). This allows comparisons of solutions from certain algorithms with the reference set to assess the performance of the algorithms.

Multi-objective Evolutionary Algorithms. Among the various strategies for solving multi-objective optimization problems, evolutionary algorithms stand out as one of the most efficient representatives to harvest a solution set in a single run while preserving diversity and convergence [29]. Starting from a random sample set drawn from the decision space, the MOEAs feature multiple major iterative loops, referred to as *generations*. A loop comprises a fixed number of evolutionary steps, called *operators*, to drive a *population* of candidate *individuals* towards the Pareto front. In the context of MOEAs, each individual corresponds to a solution that encompasses a decision vector.

Shown on the right side of Fig. 2 (B), a typical pipeline of an evolutionary algorithm starts from an initialized population $P(1) = \{p_1, p_2, \dots, p_\mu\}$ of size μ , where the i -th individual $p_i = (\mathbf{x}_{p_i}, \mathbf{f}_{p_i})$ is composed of the decision vector and the corresponding objective vector of the solution. In the iterative evolutions, the k -th generation produces a new population of solutions, denoted as $P(k)$, as the evolution outcome. The pipeline inside generation k involves the following operators:

- **Mating.** To begin with, the individuals in $P(k-1)$ are firstly divided into two sets: a mating pool $MP(k)$ and a reserved set $R(k)$. Those individuals in $MP(k)$ are then paired into various *mating pairs*:

$$M(k) = \{(p_i, p_j) \mid p_i, p_j \in MP(k), i \neq j\} \quad (2)$$

where p_i and p_j are parents for evolution in the following steps. The number of pairs is denoted as λ . Note that those *reserved* individuals in $R(k)$ will not join the crossover and mutation operations.

- **Crossover.** In this step, each mating pair (p_i, p_j) generates two offsprings by combining characteristics of the parents' decision vectors, namely, \mathbf{x}_{p_i} and \mathbf{x}_{p_j} . All the 2λ new offsprings form an offspring set

$$O(k) = \{o_i \mid 1 \leq i \leq 2\lambda\} \quad (3)$$

where $o_i = (\mathbf{x}_{o_i}, \mathbf{f}(\mathbf{x}_{o_i}))$ represents the corresponding decision and objective vectors of the offspring.

- **Mutation.** A random subset $O'(k)$ of the offspring set $O(k)$, $O'(k) \subseteq O(k)$, is selected, and all individuals in $O'(k)$ will be mutated by applying perturbation operations to the decision vectors, resulting in a mutated offspring set $\tilde{O}'(k)$. In other words, the remaining offsprings $O(k) - O'(k)$ are not mutated.

- **Environmental Selection.** To determine the optimal individual solutions, a large joint population $Q(k)$ is constructed by joining various population sets from previous steps:

$$Q(k) = R(k) \cup MP(k) \cup \tilde{O}'(k) \cup (O(k) - O'(k)) \quad (4)$$

This implies that all the parental individuals (either reserved or selected for mating) from the previous generation, along with their crossover and mutated offsprings, are subject to a competition against one another based on certain selection criteria. Such criteria usually involve measures established on individuals' objective vectors to assess their dominance and diversity [19, 29]. Individuals demonstrating a higher contribution to the overall population through optimized objective values or enhancement of the Pareto front coverage will be prioritized. Consequently, the best μ individuals that prevail in this selection process form the population $P(k)$ as the output of the current generation k .

The iterative process stops when a predefined termination criterion is satisfied, such as reaching a maximum number of generations. While MOEAs share the abovementioned pipeline for population evolution, they diverge in their specific implementations of selection, crossover, mutation, and environmental selection operators. This divergence in design choices constitutes a primary area of research and development in the evolutionary computation community.

4 DESIGN OVERVIEW

Leveraging the core principles of MOEAs, we propose a visual analytics framework designed to inspect and explain the evolutionary processes at play within these algorithms. This section outlines the requirements drawn from literature review and collaborative efforts with domain experts in evolutionary computing. Analytical tasks supported by the framework are then derived based on these established requirements.

4.1 Requirement Analysis

To accurately identify the requirements from the MOEA domain perspective, we reviewed various surveys [16, 21, 29, 49, 67] and methodological studies [1, 10, 17, 54–56] focused on evolutionary process dynamics and visualization in MOEA. This review helped identify potential research gaps. To validate and deepen our understanding,

we interviewed two evolutionary computing domain experts, E1 and E2. E1 has over a decade of research experience in MOEAs and their engineering design applications, while E2, a senior PhD student specializing in multi-objective optimization and reinforcement learning, extensively uses evolutionary algorithms. The interviews began with discussions about their MOEAs application practices and experiences with using visualizations to evaluate solution sets. We then discussed the challenges they faced with existing visualization tools and refined the open research questions identified during the literature review. The identified gaps are primarily in two areas:

- **Limited Native Visualization for Processes.** While widely-used MOEA toolkits [6, 48] incorporate a range of axis- or projection-based visualization techniques for corresponding solution sets from single generations, the process of manually plotting each generation involved in the evolutionary processes proves to be tedious. Moreover, it does not facilitate the examination of dynamics of individual solutions or the trends in populations across iterations due to difficulties in generation-wise comparison.
- **Lack of Inspection Support for Evolutionary Operators.** Another significant limitation in existing toolkits is the absence of visualizations that delve into the level of evolutionary operators. For a detailed analysis of the algorithms, it is essential to explore the mutation and selection operations of specific individuals as well as links to their corresponding parents.

Beyond the two practice gaps, we discussed challenges that hinder the development of tools to bridge these gaps. These challenges fall into three main categories. First, the massive number of individuals involved in the entire evolutionary process can lead to severe visual clutter when using standard visualization techniques. Second, juxtaposing plots from all generations does not assist in identifying temporal patterns or insights due to the typically lengthy iterative process involving hundreds or thousands of generations. The complexity increases when multiple levels of data abstraction, such as the overall statistical level [30], generation level, and individual level, are involved simultaneously. Additionally, the varied mechanisms of evolutionary operators among different algorithms make developing unified approaches to assess detailed computation processes difficult. Designing different visualizations for each algorithm can be unrealistic and impractical.

The discussions of gaps and challenges mentioned above yielded essential requirements that our framework should address, ensuring it effectively caters to the needs of researchers and practitioners in the evolutionary computation domain:

- **R1: Adopt Quality Measures as an Entrance Point.** Using aggregated quality measures as a means to comprehend the performance of MOEAs forms a fundamental basis for algorithm assessment [30]. E1 suggested that these measures could serve as high-level indicators of potential issues within the populations, thereby prompting a more in-depth analysis.
- **R2: Utilize Multiple Levels of Data Granularity.** Various levels of data granularity, including measure-level, generation-level, solution-level, and operator-level, coexist in the assessment of evolutionary processes. Integrating different levels of granularity into the same environment can facilitate a more comprehensive inspection, encompassing coarse-grained numerical indicators and fine-grained evolutionary operations.
- **R3: Locate and Connect Critical Evolution Behaviors.** Building upon the second requirement, the experts emphasized the significance of identifying representative trends and notable changes across generations. E2 noted the potential for linking observed patterns at the higher level to corresponding detailed evidence at a more granular level, enhancing the depth of analysis.
- **R4: Provide Adequate Abstraction of Operators.** In addition to the challenge of inspecting evolutionary operators, it is crucial to address the need for visualizing diverse implementations of operators. Since MOEAs can vary significantly in their operator design, particularly regarding environmental selection, developing a versatile

visualization protocol that accommodates various operator designs is essential for broad applicability.

4.2 Analytical Tasks

Building on the identified research gaps and requirements, we have distilled the following analytical tasks to guide the framework design:

T1: Understand Overall Quality Measures. Quality measures serve as the primary and most frequently-used indicators for experts and practitioners, which provide a coarse-grained overview of the distribution and performance of solutions in the objective space. This involves:

- Assessing the solution set quality in a specific generation; (**R1, R2**)
- Identifying trends within these measures as well as salient changes. (**R1, R3**)

T2: Explore Dynamics of Individuals among Generations. At a more granular level, it is essential to uncover the quality of individual solutions and their inheritance relationships across generations:

- Examining the distribution of solutions across generations and the quality of individual solutions; (**R1, R2**)
- Tracing the lineage of individuals through multiple generations. (**R3**)

T3: Inspect Detailed Evolutionary Operations. At the most detailed level of analysis, there is an interest in understanding the specific actions of evolutionary operations, including:

- Analyzing the behaviors of evolutionary operators; (**R3, R4**)
- Investigating how an individual from an older generation evolves in the subsequent generation. (**R3, R4**)

5 VISUAL ANALYTICS FRAMEWORK

Based on the identified research challenges, requirements, and analytical tasks, we introduce ParetoTracker, a visual analytics framework designed to illustrate the dynamics of population generations within evolutionary processes of MOEAs with three main components:

Performance Overview and Generation Statistics (T1). As the initial phase of the visual exploration pipeline, this component provides visual representations of algorithm performance alongside statistics of individuals across all generations, Fig. 1 (A, B). Such design enables analysts to gain a comprehensive understanding of the evolutionary process, identify trends, and discover significant patterns in the measures, thereby guiding further detailed analysis in the following components.

Visual Exploration of Individuals among Generations (T2). After identifying notable sequences of generations in the overviews, detailed information regarding decision and objective vectors is presented in the **main workspace** of the interface, Fig. 1 (C). This area facilitates lineage tracing and comparative analysis, supported by an accompanying **lineage view** shown in Fig. 1 (D).

In-depth Visual Inspection of Operators (T3). Analysts are allowed to expand the details between two consecutive generations. An **evolutionary operator view**, Fig. 1 (E), is provided to depict the intricacies of mating, crossover, mutation, and environmental selection operations on a per-individual basis, thereby offering a nuanced perspective on the evolutionary mechanisms at play.

5.1 Performance Overview and Generation Statistics

In this component, a suite of quality measures, along with corresponding visualization modules, is utilized to enable a comprehensive analysis of the generations (**T1, R1**). Additionally, statistics regarding the type of evolutionary operations are provided to offer insights into the evolutionary processes occurring in each generation.

Generation-level Quality Measures. Drawing from established literature on the quality assessment of solution sets in multi-objective optimization [30], the following four quality measures have been selected to evaluate generations. These measures fall into two categories: aggregated measures, which assess multiple aspects of quality simultaneously, and specialized measures, which focus on a specific aspect.

- *Inverted Generational Distance (IGD) and Hypervolume (HV)*: IGD and HV are among the most frequently-used aggregated measures in the MOEA literature by simultaneously evaluating both convergence and diversity. IGD calculates the average distance from the solutions in the population to the points in the reference set in the objective space, whereas HV quantifies the volume of the region dominated by the population relative to a predefined point in the objective space. A lower IGD or a higher HV value signifies a population with superior convergence and diversity.
- *Spacing (SP) and Maximum Spread (MS)*: In contrast to the aggregated quality indicators, SP and MS focus on specific aspects of the quality of solution sets. SP computes the uniformity of the solutions by measuring the variance in distances between them, with a larger SP value indicating a less uniform distribution of solutions. MS evaluates the extent of the solutions across each objective, with a larger MS value denoting a broader spread of solutions.

As depicted in Fig. 1 (A), the selected quality measures are visualized using four separate quality line charts in the quality line chart view, each corresponding to one measure. The horizontal axis, aligned across all line charts, denotes the sequence of generations, while the vertical axis represents the corresponding measure values. A semantic zoom feature is implemented on the horizontal axis for detailed examination of specific generation ranges, and the zoom ratios are synchronized to support alignment and comparison across all line charts.

Statistics of Populations. Understanding the origins of individuals concerning evolutionary operations is vital for analyzing population dynamics. Specifically, analysts need to discern whether an individual in a given generation is a directly inherited parent from the previous generation or a mutated offspring evolved from a parental pair. To address this, we employ a stacked bar design visualization to depict individual proportions based on their origins. As shown in Fig. 1 (B), the ordinal indices of generations are placed on the stacked bars' left side. Inside the bar for the k -th generation, a series of segments represent the proportions of the following four types of origins:

- 1) Individuals reserved from the previous generation, $R(k)$;
- 2) Individuals in the mating pool, $MP(k)$;
- 3) Crossover offsprings based on their parental pairs, $O(k) - O'(k)$;
- 4) Mutated offspring, $\tilde{O}'(k)$.

A categorical color scheme is employed to facilitate distinction among the four origins, where light blue is for **reserved individuals**, green for **individuals from the mating pool**, purple for **crossover offspring**, and yellow for **mutated offspring**. It is worth-noting that this color mapping is consistently applied in the entire framework. Additionally, bar segments featuring striped patterns indicate the proportions of individuals that did not survive the environmental selection process.

5.2 Population Analysis of Generation Sequences

Delving into the specifics of solutions, including decision and objective vectors of individuals across all generations, facilitates a granular understanding of evolutionary dynamics beyond mere aggregate measures and statistics (**T2**, **R3**). This component leverages an abstraction of individuals and their lineages, displayed in the main workspace of the interface, alongside a lineage view, Fig. 1 (C, D).

Decision and Objective Space Analysis. When notable patterns like salient increases or decreases in quality measure values are discovered, analysts can select a range of generations by brushing along the horizontal axis or clicking on any data point in quality line charts. The detailed information for the selected generations is then depicted in a juxtaposed manner in the main workspace, Fig. 1 (C). Following **R4** for providing appropriate data abstraction at each analytical level, solutions are abstracted as two sets of vectors, i.e., decision and objective vectors, and are visualized in two vertically-stacked scatterplots, with the top one for the objective space and the bottom one the decision space.

Layout: It is common for the dimensions of decision and objective spaces to exceed two, sometimes reaching beyond ten. To address this, projection techniques are employed, utilizing PCA for the objective space and t-SNE [52] for the decision space. To ensure comparability, we address the alignment of projection results within the same space:

- For the objective space, the PCA projection matrix is fit using the reference set and subsequently applied across all scatterplots, enabling clear observation of solution distributions relative to the ground-truth reference points.
- For the decision space, t-SNE parameters are trained based on the union of all decision vectors present in the selected generations, with these parameters consistently applied across all selected generations.

It should be noted that the choice of the projection techniques is based on the typical configuration in multi-objective optimization problems, where the count of objective functions generally does not exceed 10, often hovering around 4. Conversely, the dimensionality of decision space can span from tens to thousands [49]. Given this disparity, linear projection techniques may fall short in adequately unveiling the complex distributions and patterns in the decision space.

Visual Encoding: The dot colors are mapped to the origins of evolutionary operations based on the abovementioned categorical color scheme in Sec. 5.1. To highlight the quality of individuals, analysts can assign dot sizes to the values of two individual-level quality measures: 1) the distance to the nearest reference point [53], and 2) the distance to the nearest solution in the objective space. The former size mapping serves as a visual cue for the convergence of solutions towards the Pareto front, i.e., the accuracy of the solutions, and the latter one illustrates the uniformity of distributions, with smaller dot sizes implying shorter distances. If clusters of small-sized dots are observed, it strongly suggests that the corresponding regions may contain a considerable number of duplicated solutions. On the other hand, large and scattered dots indicate regions that are relatively sparse. Besides the objective space, the decision space scatterplots incorporate an additional size mapping mode by utilizing the distance to the nearest solution within the decision space, which further reveals the spatial relationships among solutions in the decision space. To illustrate the outcomes of the environmental selection process for each generation, individuals that do not survive are distinctly marked as crosses instead of dots in the scatterplots and will not appear in successive generations.

For test problems with reference sets, a straightforward visualization approach is to plot the corresponding objective vectors of reference points in the objective space scatterplot. However, this can lead to significant visual clutter due to the high number of sampled reference points typically involved [19, 22]. To address this, we adopt a density map to illustrate the distribution of the reference points. The density map is positioned beneath all marks representing individuals to provide an unobstructed view, effectively enabling comparisons between the current population spread and the desired Pareto front.

Interaction: When hovering the mouse pointer over a dot, a tooltip can be activated to display critical attributes about the solution. The information includes the generation order, individual ID, the origin of evolutionary operator, the distance to the nearest reference point distance, the values on all objectives, and the survival status after the environmental selection process in the corresponding generation.

Lineage Analysis. Beyond simply placing scatterplots of all selected generations side by side, our framework also facilitates a detailed depiction of the parental relationships between generations. Selecting any dot in a scatterplot triggers the display of connecting curves that link the individual either to its parents (or back to itself if it is a direct inheritance from the preceding generation), Fig. 1 (C). These parental links extend all the way to the initial generation in the chosen range. The color of the curves corresponds to the offspring's origin of evolutionary operation, while the curve thickness illustrates the objective space distance between the offspring and the linked parent. Enable the lineage connections for multiple individuals simultaneously is also supported.

Analysts might be curious about a group of individuals to determine their long-term evolutions as well as whether they evolved from a common set of ancestors. However, visualizing lineage trees in scatterplots across a considerable amount of generations may lead to visual clutter, especially when comparing multiple trees for several individuals. To address this, a dedicated *lineage view* is designed to clarify the illustration of multiple lineages, Fig. 1 (D) and Fig. 3, with two panels:




Fig. 3: Lineage view. (Left Panel) A timeline of the selected individuals. (Right Panel) A scatterplot depicting the orders of ancestors with arrows.

Left Panel: The horizontal axis represents the sequential order of generations, with each lineage-enabled individual represented as a timeline in a row. The timeline highlights the originating generation of the individual and the point at which it no longer passes the environmental selection process. For illustrating common ancestors among individuals, cross-row curved links are used for showing the generation where a closest common ancestor is identified. Note that the range of generations covered by the horizontal axis only spans the necessary range of the selected individuals instead of depicting all generations in the evolutionary process. The axis starts from the earliest generation where a common ancestor appears or an individual is born, and it ends with the final survived generation among all selected individuals.

Right Panel: In addition to the temporal comparisons, this panel visualizes the distribution of the selected individual and its lineage across the generations on the two scatterplots for decision and objective vectors, respectively. Arrows pointing from parents to offsprings are rendered in gradient shades where lighter tones indicate earlier generations. This assist in examining the decision and objective vector distributions of all related individuals on the lineage tree collectively.

5.3 In-depth Inspection of Evolutionary Operations

At the most detailed level of analysis, this component offers insights into the specific operations executed on the population within one generation to create the subsequent generation. When a range of generations is selected in the main workspace, clicking on the “plus” sign between the scatterplots of two consecutive generations unfolds an *evolutionary operator view* which shows the intricate steps of each evolutionary operation. Drawing from the foundational principles of evolutionary computing outlined in Sec. 3, we specify these operations within the multi-objective optimization context, aiming to provide data abstraction that aligns with analysts’ requirements as suggested in R4.

Data Abstraction for Operators in MOEAs. Through discussions with domain experts and an examination of prevalent MOEA tools like PlatEMO [48] and Pymoo [6], we have identified detailed abstractions for mating, crossover, mutation, and environmental selection operations, encompassing a broad spectrum of MOEAs.

Mating and Crossover: The common mating protocol used in MOEAs is the same as how $M(k)$ is constructed in Sec. 3. For crossover, we focus on the Simulated Binary Crossover (SBX) operator [12], tailored for real-valued recombination operations. Given a pair of individuals (p_i, p_j) with their decision vectors \mathbf{x}_{p_i} and \mathbf{x}_{p_j} , SBX yields two offsprings o_i and o_j with decision vectors calculated as follows:

$$\mathbf{x}_{o_i} = 0.5 \left[(1 + \beta) \mathbf{x}_{p_i} + (1 - \beta) \mathbf{x}_{p_j} \right] \quad (5)$$

$$\mathbf{x}_{o_j} = 0.5 \left[(1 - \beta) \mathbf{x}_{p_i} + (1 + \beta) \mathbf{x}_{p_j} \right] \quad (6)$$

where β is a spread factor derived from a specific probability distribution. In a more intuitive explanation, this approach linearly combines the parents’ decision vectors with symmetric random perturbations, and all four vectors are along the same line in decision space. Notably, some implementations [6] introduce an additional post-crossover random perturbation to enhance the genetic diversity of the offsprings.

Mutation: Alongside SBX, polynomial mutation is frequently employed. For each selected offspring in $O'(k)$ undergoing mutation, a random perturbation vector $\delta \in \mathbb{R}^n$ is generated in a defined range to avoid excessive and unrealistic modifications.

Environmental Selection: At a high level, this process involves labeling candidate individuals as “survived” or “died”. Yet, this abstraction




Fig. 4: Examples of the glyphs for (a) the mating and crossover panel and (b) the mutation panel.

may not fully capture the nuanced selection mechanisms inherent in MOEAs. After exploring various search and selection strategies of algorithms [60], we suggest a more intricate two-level abstraction: grouping and fitness scoring.

- Candidates are initially sorted into groups based on specific criteria. Strategies like non-dominated sorting [11] prioritize groups so that individuals in lower-priority groups are dominated by at least one individual in higher tiers. Conversely, methods like uniform space partitioning [9] treat all groups as equal in priority.
- Within these groups, individuals receive a fitness score that dictates their competitive standing for survival. As mentioned in Sec. 3 on environmental selection operators, such a fitness score reflects an individual’s contribution to the population in terms of its proximity to the Pareto front and its coverage on the Pareto front. For instance, NSGA-II [11] uses crowding distances to eliminate duplicated solutions in densely populated regions while favoring solutions in less crowded areas, thereby enhancing population diversity. MOEA/D [64] ranks solutions using a weighted-sum method over objective values in their corresponding objective vectors. However, the exact method for computing fitness scores can differ across various MOEAs, leading to our design choice of using numerical scores rather than illustrating the detailed computation process. Note that in the prioritized grouping settings, the primary determinant is group rank, followed by individual fitness scores. For non-prioritized groupings, quotas for survival are assigned in each group based on fitness scores.

Visualization and Interaction. The evolutionary operator view is divided into three distinct panels, each corresponding to one of the evolutionary operations previously discussed. The basic layout for these panels includes a list of individuals alongside a pair of scatterplots for the decision and objective spaces. However, each panel features a tailored design to effectively represent the details of different operators.

Mating and Crossover Panel: This panel displays each pair of parents from the previous generation and their offsprings in a list format where a glyph design, Fig. 4, is employed to illustrate the relationships between parents and offsprings. In the glyph, two green dots represent the parents, and two tiny purple dots represent the offspring. The connecting line’s length between the two parent dots signifies the Euclidean distance between the parent individuals in the decision space. A Gaussian-blurred halo around the purple offspring dots indicates the degree of post-hoc random perturbations applied to the offspring before mutation, with a larger halo showing a greater perturbation. Dashed concentric circles around the dots denote the quality of the individuals, indicating whether the crossover operation has produced offsprings superior to the parents. The radii of these circles are proportional to the distance to the nearest reference point, which is the same individual-level quality measure utilized in Sec. 5.2. The scatterplots for decision and objective spaces display the spatial distributions of the parents, the reserved individuals from the previous generation, and the offspring within these two spaces. The rows in the list can be sorted by parent-parent or parent-offspring distances.

Mutation Panel: A similar glyph design is used to list all mutated offsprings, Fig. 4 (b), where purple and yellow dots represent the crossover offsprings and its mutated counterparts, respectively. The linkage length between dots illustrates the Euclidean distance the offspring has mutated in the decision space. The same dashed concentric circles around the dots convey the individual’s quality, specifically its convergence to the Pareto front. To illustrate the direction of mutations in the decision space, a bar chart is utilized for each mutation to




Fig. 5: The evolutionary operator view for the 59th generation.

display the values of all dimensions in the mutation direction vector, represented as $\mathbf{x}_{\text{mutation}} - \mathbf{x}_{\text{offspring}}$ in \mathbb{R}^n . The values along the same dimension are subjected to min-max normalization to $[0, 1]$, facilitating a straightforward comparison across different mutations and dimensions. List sorting is supported based on different criteria, including distances between offsprings and corresponding mutated results as well as dimensions of the mutation direction vectors.

Environmental Selection Panel: Individuals entering the environmental selection process are initially sorted into groups, with each group displayed in the list according to their priority. Individual IDs, evolutionary operation origins (indicated by the background color), and fitness scores are depicted in small rectangles within each group. The individuals in groups are further organized by fitness scores, highlighting the top-performing individual. The origin proportions of evolutionary operations within each group are represented as stacked bars, consistent with the design for generation statistics in Sec. 5.1. The total length of the stack bars across all groups is normalized based on the number of individuals in the groups, with the total and survived individual counts displayed at the group's header.

Interactions. The evolutionary operator view facilitates rich interactions. Clicking on any list row or scatterplot dot highlights corresponding individuals across the entire view; in addition, linking lines between dots are displayed simultaneously to visualize relationships, including parental links in the mating and crossover panel and the connection between offspring and their mutations. Toggle button in the group headers enables simultaneous highlighting of all individuals in a group. The tooltip feature for dots of individuals, as discussed in Sec. 5.2, is also applied across the scatterplots in the panels.

6 EVALUATION

In this section, we discuss the results from case studies and expert interviews to demonstrate the effectiveness of ParetoTracker on benchmarking and real-world multi-objective optimization problems. ParetoTracker is implemented with a browser-server architecture by employing Python Flask for the server-side logic as well as Vue3 and D3.js for the client side. Pymoo [6] is used for running MOEAs on test problems.

6.1 Case Study 1: SMS-EMOA on DDMOP2 Problem

In the first case study, we employ the DDMOP2 test problem from the DDMOP test suite [18], where the problems are devised from scenarios encountered in real-world applications. The DDMOP2 problem comprises five decision variables which represent the reinforcing components of automobile frontal structures while aiming to minimize three objectives. The SMS-EMOA algorithm [5] is selected for this task due

to its prominence among indicator-based MOEAs. SMS-EMOA prioritizes individuals based on their contribution to the Hypervolume (HV) measure, which is indicative of an individual's importance in the environmental selection process, and such approach is known for ensuring good coverage and uniformity along the Pareto front. The population size is set to 100, and the algorithm run lasts 500 generations.

Quality Measures and Generation Statistics (T1). After the data is loaded, an initial review of quality line charts and generation stacked bars provides an overview of the evolutionary process, Fig. 1 (a.1). The two measures, HV and IGD, show relatively monotonic trends towards their desired optimality, which is characterized by increasing HV and decreasing IGD values. Meanwhile, SP and MS display fluctuations in the early evolutionary stages (up to the 40th generation) before stabilizing. Typically, changes in SP and MS are strong indicators that the coverage and spread of the population on the Pareto front are still expanding, a phenomenon primarily triggered by mutated individuals. Following the 40th generation, these measures begin to stabilize, suggesting that little to no effective mutations are occurring, and the solution distribution is starting to reach a steady state. Such stabilized status can be confirmed by the stacked bars, which show a decreasing number of survived mutated individuals after the 40th generation, Fig. 1 (b.1, b.2). Notably, in the quality line charts, a small increase in MS is observed between the 58th and 87th generations, accompanied by a similar value rise in SP, indicating noteworthy variations in solution distributions that requires closer examination, Fig. 1 (a.2).

Generation-level Analysis (T2). Triggered by the notable increase in MS and SP, a detailed examination of solution distributions is conducted by activating the scatterplot series around the 58th generation. After clicking on the point representing the 58th generation in the MS line chart, scatterplots for three consecutive generations both prior to and following the 58th generation can be displayed in the main workspace, Fig. 1 (C). In the results, an important observation between the 58th and 59th generations is the emergence of a mutated individual (ID #5856) in the 59th generation, which broadens the solution set coverage in the top left corner of the objective space scatterplot relative to the blue density map of reference points, Fig. 1 (c.1). This mutation contributes to the MS increase by expanding the range of the solution set with respect to values covered in the objective space. Additionally, by employing the dot size mapping based on nearest neighbor distances in the objective space, it can be observed that the dot size of #5856 is significantly larger than most other individuals, suggesting its role in contributing to the increase of SP value observed in the line chart.

Upon further investigation by clicking on the dot of #5856, lineage connections displayed in the preceding scatterplots reveal that #5856

emerged from a mutation involving a pair of parents, one of which, #5680 from the 58th generation, is extremely close to the offspring #5856 in the objective space (Fig. 1 (c.2)). The timeline in the lineage view indicates that #5856 does not persist beyond the 71st generation, Fig. 1 (d.1). Nonetheless, the MS and SP measures remain elevated until the 87th generation. This observation implies that other solutions in the vicinity of #5856 may have taken over its role in the solution space. This hypothesis is validated by examining scatterplots around the 71st generation, where another individual, #7054, appears in nearly the identical location as #5856, Fig. 1 (d.2, d.3).

Evolutionary Details between Generations (T3). To conduct a fine-grained analysis, we delve into the specifics of how #5856 emerged in the 59th generation and the circumstances leading to the demise of #7054 in the 87th generation. By activating the evolutionary operator view for the 59th generation and selecting #5856 in the objective scatterplot in the environmental selection panel, all elements related to #5856 are highlighted for analysis, Fig. 5.

In the mating and crossover panel, the offspring from the crossover operation does not initially occupy the position of #5856, Fig. 5 (a). Instead, it is the subsequent mutation operation that propels the offspring towards the top left part of the cluster, thus shaping the distinctive characteristics of #5856, Fig. 5 (b). In the grouped list of the environmental selection panel, #5856 ranks second in fitness score in the first group, Fig. 5 (c), reflecting its significant contribution to the HV indicator in SMS-EMOA and, by extension, to the overall solution set quality. This observation extends to other mutated individuals like #5877 and #5849, Fig. 5 (d), who exhibit a similar expansion in solution set coverage at critical cornering locations. Such patterns suggest a possible algorithmic preference for individuals enhancing Pareto front coverage.

Turning to evolutionary operator view of the 87th generation, Fig. 1 (e.1), the environmental selection panel reveals that #7054 falls in the second group where all members did not survive. This situation results from the non-dominated sorting mechanism in SMS-EMOA, where #7054 is outperformed and thus dominated by at least one individual from the prioritized first group. A closer inspection of objective vector values in Fig. 1 (e.2) indicates that #7054 is dominated on all objectives by a mutated individual, #8605, leading to its elimination from the evolutionary process. This analysis reveals the dynamic interplay of evolutionary operators and selection pressures that shape the development and eventual pruning of solutions in MOEAs.

6.2 Case Study 2: NSGA-II on DTLZ3 Problem

The second case study explores the DTLZ3 problem from the widely-used DTLZ benchmark suite for multi-objective optimization [13], which features 10 decision variables and 3 objectives. The NSGA-II [11] algorithm is selected for this analysis. We adopt the same population and generation settings as in the first case study.

Quality Measures and Generation Statistics (T1). Observations from the quality line chart and generation stacked bars, Fig. 6 (A), highlight a notable uptick in the HV measure at the 371st generation, alongside an accompanying drop in IGD pointing towards an optimal trend. This pattern strongly suggests that the population may have begun converging on the Pareto front after an extensive exploration phase where the measures do not vary too much; prior to the population converging on the Pareto front, it is likely that it remains considerably distant from the target Pareto front.

Generation-level Analysis (T2). The notable shifts observed in the HV and IGD metrics necessitate a closer examination of the scatterplot series around the 371st generation, Fig. 6 (B.1). Prior to the 371st generation, a cluster of individuals is situated near the bottom-left part of the reference point density map. A critical transition occurs in the 371st generation, where the majority of these individuals do not survive the environmental selection process, leaving behind only one mutated survivor, #37085. From the 372nd generation onwards, the emergence of new crossover and mutated individuals is noted in this region.

When switching the dot size mapping to the nearest reference point distances in the enabled scatterplots, Fig. 6 (B.2), an abundance of yellow crosses in the decision space scatterplots have appeared, with their

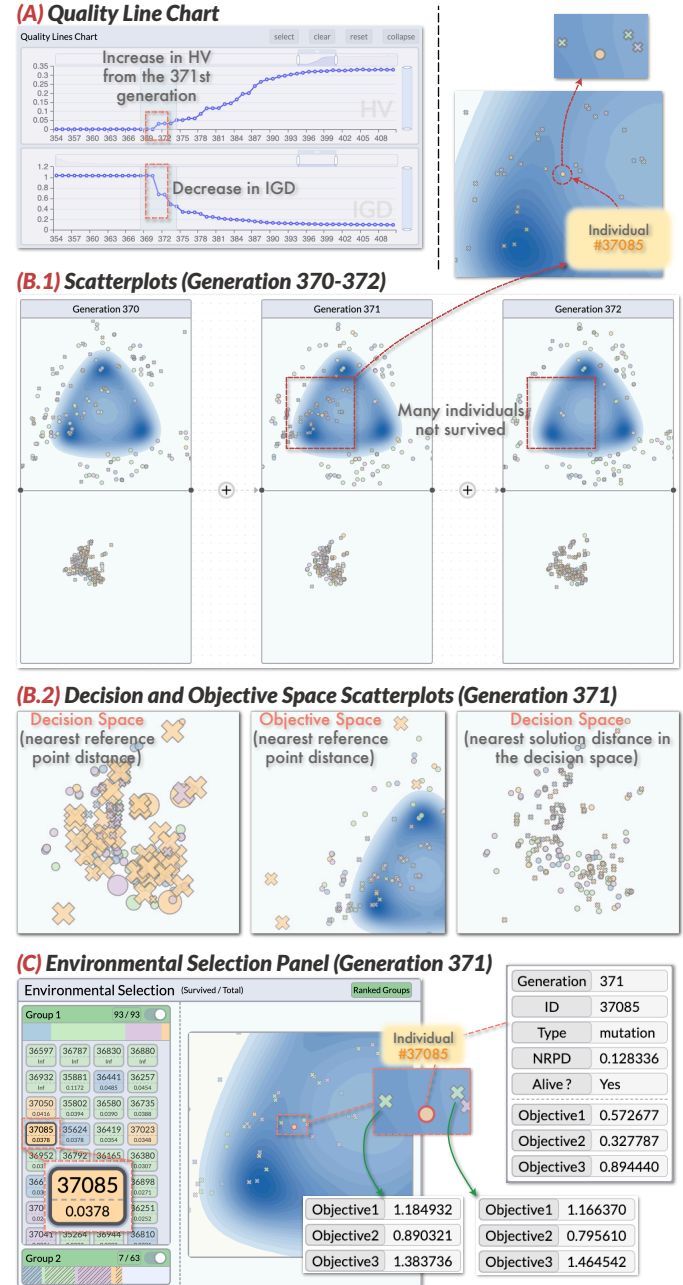


Fig. 6: The result of NSGA-II on DTLZ3. (A) The HV measure starts to increase from the 371st generation. (B.1) Individual #37085 in the 371st generation has dominated a considerable amount of nearby individuals. (B.2) The distributions of decision and objective vectors exhibit a large incoherence. (C) The details of how #37085 dominates other individuals are illustrated.

counterparts rarely visible in the objective space scatterplots. Panning the viewport of the objective space scatterplots uncovers these crosses positioned far from the reference density map area. This observed phenomenon contrasts sharply with the distribution patterns in the t-SNE projection results in the decision space scatterplots, where individuals appear to be relatively uniformly distributed. Such uniformity can further be confirmed when applying dot size mapping to the nearest solution distances in the decision space, where the dot sizes have very few variations. The discrepancy between the uniform distribution in the decision space and the varied distribution in the objective space underscores the high sensitivity of the objective functions to the input decision vectors. In other words, small changes in the decision space can lead to significant differences in the objective space, which suggests high complexity of the decision space landscape [40].

Evolutionary Details between Generations (T3). Further examination is directed towards #37085 in the 371st generation to understand its survival, Fig. 6 (C). Highlighting #37085 in the environmental selection panel for this generation reveals its placement in the first group based on the non-dominated sorting mechanism in NSGA-II. Examination of the nearest reference point distances via tooltips shows that #37085 exhibits significantly lower values, with its objective vector values also lower across all objectives compared to surrounding non-surviving individuals. These insights confirm #37085’s superiority over its nearby counterparts, ensuring its preservation within the population.

6.3 Expert Interview

The visual analytics framework was further evaluated through interviews with four domain experts, including those two from the preliminary interviews (E1 and E2) and two additional experts, denoted as E3 and E4, respectively. E3, a senior researcher, has extensive experience in developing optimization algorithms for real-world applications. E4, a graduate student, concentrates on multi-objective optimization research. The interview process began with an introduction to the framework and a demonstration of its analytical features. Then, we encouraged the experts to freely explore the system using pre-loaded data corresponding to the two case studies. They were also guided to go through the analysis processes that led to the results presented in the case studies. Feedback was collected on findings in case studies, functionality for data analysis, and visualization design of the framework. Interviews lasted between 30 minutes to 1.5 hours.

Experts’ Comments on Case Studies. The experts provided several pieces of feedback regarding the observed evolutionary processes in the case studies. For the DDMOP2 problem, where a mutated individual #5856 initiated an exploration in the objective space to enhance the population’s coverage but ultimately failed, experts posited two possible explanations. First, the algorithm’s characteristics may align perfectly with the nature of the test problem, leading to the population converging to the Pareto front around the 40th generation and stabilizing. In this scenario, mutations like #5856 are unnecessary. However, the experts emphasized that this conclusion should be assessed by comparing runs from different algorithms to further evaluate the convergence. Second, the algorithm may discourage mutations that increase diversity and consistently generate individuals with higher accuracy, even if they are duplicates in certain local areas. The experts suggested that improvements could be made by adjusting environmental selection strategies in the next step to see results under different mutation tolerances.

Conversely, mutation #37085 in the DTLZ3 problem illustrates a successful exploration towards the Pareto front, which results in the ultimate success of convergence despite being continually poor according to the quality line charts (starting to converge only after the 371st generation). The experts noted that it could be an interesting topic to further investigate the underlying factors that generate successful mutations, as this could inform future algorithm design.

Overall Analysis Pipeline. The experts expressed enthusiasm for the incorporation of interactive visualization techniques to examine evolutionary processes in MOEAs. They commended the capacity of the framework to represent generations and individuals across various aspects, including quality measures, generation statistics, and detailed process views. E1 highlighted the advantage of ParetoTracker’s multi-faceted design over traditional systems that rely on static visualization, noting its facilitation of interactive exploration and pattern discovery across generations. E4 recognized the potential of ParetoTracker for integrating into daily research workflows, praising its adaptability for comprehensive analysis of different algorithms. “The proposed data abstraction method accommodates a majority of MOEA algorithms. By adapting the logging mechanisms of evolutionary computing processes, ParetoTracker could be seamlessly connected to computational tools once the output meets the data protocol,” remarked E4.

Visualization and Interaction. The interactive visualization design received positive feedback for its efficacy in depicting various measures, solution distributions, and the environmental selection process across generations. E2 mentioned that patterns identified at higher analytical

levels could serve as cues for subsequent detailed investigations, thereby enriching the insights obtained. E3 emphasized the value of a detailed view of evolutionary operations, stating, “Sometimes the algorithm gets stuck in suboptimal conditions as indicated by quality measures. The individual-level visualization, which illustrates the creation and disappearance of individuals, can help identify pivotal evolutionary events that either enhance or hinder optimization efforts.”

Suggestions on Improving the Current Framework. The experts offered suggestions for potential future improvements to the framework. E4 suggested extending the data abstraction protocol for evolutionary operations to include mechanisms for introducing new, randomly selected individuals into the population, thereby increasing diversity. E3 recommended incorporating contextual information or visualizations related to the specific application domains of test problems, which could enhance the applicability of the framework in particular scenarios.

7 DISCUSSION AND CONCLUSION

We present ParetoTracker, a visual analytics framework illustrating the evolutionary dynamics in multi-objective evolutionary algorithms. Its multi-level design allows analysts to explore evolutionary processes from different angles, including quality measures, generation statistics, solution distributions across generations, and evolutionary operator actions. This holistic approach enables analysts to go beyond conventional single solution set analysis, promoting a deeper exploration into the evolutionary operations of each generation.

Comparison with Existing Tools. Our framework is compared with existing MOEA frameworks and toolkits [6, 48], focusing on evolutionary dynamics analysis support. These major computation toolkits offer static visualization for quality measures and solution sets through non-interactive charts and scatterplots, as identified in the research gaps in Sec. 4.1. ParetoTracker enhances these functionalities by providing an integrated visual analytics environment with detailed visualizations of generation sequence and evolutionary operators. This facilitates more effective inspection of individuals and operators across generations, a feature not addressed by the computation toolkits. With the additional design of rich interactions among multiple views, ParetoTracker enables multi-level analysis from quality measures to operators, which assists in identifying notable patterns in evolutionary processes and correlating pattern influences at various granularity.

Scalability. Our current framework effectively manages data logging for standard experimental setups, handling population sizes per generation up to several hundred. Yet, in large-scale optimization scenarios [49] with significantly larger dimensionality and population sizes, the log data volume for evolutionary processes can quickly grow. We plan to incorporate advanced sampling strategies prioritizing critical generation subsequences in future implementations. Additionally, the risk of visual clutter in scatterplot-based visualizations requires scatterplot simplification methods [44, 63] to address this problem.

Generalizability. The data abstraction framework designed is adaptable to most MOEAs by offering ample detail and ensuring adaptability for various environmental selection processes. However, applying it to decomposition-based MOEAs is challenging due to their complex selection mechanisms that resist straightforward abstraction. For these MOEAs, our current strategy uses a higher-level abstraction that only tracks individuals’ survival status. Future work will focus on creating more specialized methods to better support such MOEAs.

Limitations and Future Work. Future enhancements could include a plugin-based environment for easy integration of new evolutionary operators. Additionally, incorporating application-specific information could improve utility, particularly in multiple-criteria decision-making applications without reference points, enabling relative comparison between generations and different algorithm runs. This would allow analysis without ground truth, enhancing the framework’s applicability and real-world value. To extend to a broader audience, including educational settings, we plan to assess ParetoTracker’s effectiveness across diverse user groups with varying expertise levels.

SUPPLEMENTAL MATERIALS

The supplemental materials include a demo video of our framework and the corresponding subtitle file of the video. An implementation is released at <https://github.com/VIS-SUSTech/ParetoTracker>.

ACKNOWLEDGMENTS

This work was supported in part by the National Natural Science Foundation of China (No. 62202217), Guangdong Basic and Applied Basic Research Foundation (No. 2023A1515012889), Guangdong Key Program (No. 2021QN02X794), and Guangdong Natural Science Funds for Distinguished Young Scholar (No. 2024B1515020019).

REFERENCES

- [1] I. Alberto, C. A. Coello Coello, and P. M. Mateo. A comparative study of variation operators used for evolutionary multi-objective optimization. *Information Sciences*, 273:33–48, 2014. doi: 10.1016/j.ins.2014.03.042 3
- [2] N. Andrienko, G. Andrienko, S. Miksch, H. Schumann, and S. Wrobel. A theoretical model for pattern discovery in visual analytics. *Visual Informatics*, 5(1):23–42, 2021. doi: 10.1016/j.visinf.2020.12.002 2
- [3] R. Arbolino, R. Boffardi, L. De Simone, and G. Ioppolo. Multi-objective optimization technique: A novel approach in tourism sustainability planning. *Journal of Environmental Management*, 285:112016, 2021. doi: 10.1016/j.jenvman.2021.112016 1
- [4] R. A. Becker and W. S. Cleveland. Brushing scatterplots. *Technometrics*, 29(2):127–142, 1987. doi: 10.1080/00401706.1987.10488204 2
- [5] N. Beume, B. Naujoks, and M. Emmerich. SMS-EMOA: Multiobjective selection based on dominated hypervolume. *European Journal of Operational Research*, 181(3):1653–1669, 2007. doi: 10.1016/j.ejor.2006.08.008 7
- [6] J. Blank and K. Deb. Pymoo: Multi-objective optimization in python. *IEEE Access*, 8:89497–89509, 2020. doi: 10.1109/ACCESS.2020.2990567 4, 6, 7, 9
- [7] J. Chae, C. Steed, and G. D. Tourassi. Visualization for classification in deep neural networks. In *Proceedings of the Workshop on Visual Analytics for Deep Learning*, 2017. 2
- [8] A. Chatzimparmpas, R. M. Martins, K. Kucher, and A. Kerren. VisEvol: Visual analytics to support hyperparameter search through evolutionary optimization. *Computer Graphics Forum*, 40(3):201–214, 2021. doi: 10.1111/cgf.14300 2
- [9] R. Cheng, Y. Jin, M. Olhofer, and B. Sendhoff. A reference vector guided evolutionary algorithm for many-objective optimization. *IEEE Transactions on Evolutionary Computation*, 20(5):773–791, 2016. doi: 10.1109/TEVC.2016.2519378 6
- [10] A. De Lorenzo, E. Medvet, T. Tušar, and A. Bartoli. An analysis of dimensionality reduction techniques for visualizing evolution. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 1864–1872, 2019. doi: 10.1145/3319619.3326868 2, 3
- [11] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002. doi: 10.1109/4235.996017 6, 8
- [12] K. Deb, K. Sindhya, and T. Okabe. Self-adaptive simulated binary crossover for real-parameter optimization. In *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 1187–1194, 2007. doi: 10.1145/1276958.1277190 6
- [13] K. Deb, L. Thiele, M. Laumanns, and E. Zitzler. Scalable multi-objective optimization test problems. In *Proceedings of the Congress on Evolutionary Computation*, vol. 1, pp. 825–830, 2002. doi: 10.1109/CEC.2002.1007032 8
- [14] A. Delaforge, J. Aze, S. Bringay, C. Mollevi, A. Sallaberry, and M. Servajean. EBBe-Text: Explaining neural networks by exploring text classification decision boundaries. *IEEE Transactions on Visualization and Computer Graphics*, 29(10):4154–4171, 2023. doi: 10.1109/TVCG.2022.3184247 2
- [15] A. Ender, W. Ribarsky, C. Turkay, B. W. Wong, I. Nabney, I. D. Blanco, and F. Rossi. The state of the art in integrating machine learning into visual analytics. *Computer Graphics Forum*, 36(8):458–486, 2017. doi: 10.1111/cgf.13092 2
- [16] J. G. Falcón-Cardona and C. A. Coello. Indicator-based multi-objective evolutionary algorithms: A comprehensive survey. *ACM Computing Surveys*, 53(2), 35 pages, 2020. doi: 10.1145/3376916 3
- [17] B. Filipič and T. Tušar. A taxonomy of methods for visualizing pareto front approximations. *Proceedings of the Genetic and Evolutionary Computation Conference*, pp. 649–656, 2018. doi: 10.1145/3205455.3205607 2, 3
- [18] C. He, Y. Tian, H. Wang, and Y. Jin. A repository of real-world datasets for data-driven evolutionary multiobjective optimization. *Complex and Intelligent Systems*, 6(1):189–197, 2020. doi: 10.1007/s40747-019-00126-2 7
- [19] Z. He and G. G. Yen. Visualization and performance metric in many-objective optimization. *IEEE Transactions on Evolutionary Computation*, 20(3):386–402, 2016. doi: 10.1109/TEVC.2015.2472283 2, 3, 5
- [20] F. M. Hohman, M. Kahng, R. Pienta, and D. H. Chau. Visual analytics in deep learning: An interrogative survey for the next frontiers. *IEEE Transactions on Visualization and Computer Graphics*, 25(8):2674–2693, 2018. doi: 10.1109/TVCG.2018.2843369 2
- [21] Y. Hua, Q. Liu, K. Hao, and Y. Jin. A survey of evolutionary algorithms for multi-objective optimization problems with irregular pareto fronts. *IEEE/CAA Journal of Automatica Sinica*, 8(2):303–318, 2021. doi: 10.1109/JAS.2021.1003817 3
- [22] Y. Huang, Z. Zhang, A. Jiao, Y. Ma, and R. Cheng. A comparative visual analytics framework for evaluating evolutionary processes in multi-objective optimization. *IEEE Transactions on Visualization and Computer Graphics*, 30(1):661–671, 2024. doi: 10.1109/TVCG.2023.3326921 2, 5
- [23] A. Ibrahim, S. Rahnamayan, M. V. Martin, and K. Deb. 3D-RadVis: Visualization of Pareto front in many-objective optimization. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pp. 736–745, 2016. doi: 10.1109/CEC.2016.7743865 2
- [24] A. Inselberg. *Parallel Coordinates: Visual Multidimensional Geometry and Its Applications*. Springer New York, 2009. doi: 10.1007/978-0-387-68628-8 2
- [25] I. Jolliffe. *Principal Component Analysis*. John Wiley & Sons, Ltd, 2014. doi: 10.1002/9781118445112.stat06472 2
- [26] M. Kahng, N. Thorat, D. Horng, P. Chau, F. B. Vi, and M. Wattenberg. GAN Lab: Understanding complex deep generative models using interactive visual experimentation. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):310–320, 2019. doi: 10.1109/TVCG.2018.2864500 2
- [27] H. Kim, J. Choo, C. Lee, H. Lee, C. K. Reddy, and H. Park. PIVE: Per-iteration visualization environment for real-time interactions with dimension reduction and clustering. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 1001–1009, 2017. doi: 10.1609/aaai.v31i1.10628 2
- [28] B. La Rosa, G. Blasilli, R. Bourqui, D. Auber, G. Santucci, R. Capobianco, E. Bertini, R. Giot, and M. Angelini. State of the art of visual analytics for explainable deep learning. *Computer Graphics Forum*, 42(1):319–355, 2023. doi: 10.1111/cgf.14733 2
- [29] K. Li, R. Wang, T. Zhang, and H. Ishibuchi. Evolutionary many-objective optimization: A comparative study of the state-of-the-art. *IEEE Access*, 6:26194–26214, 2018. doi: 10.1109/ACCESS.2018.2832181 3
- [30] M. Li and X. Yao. Quality evaluation of solution sets in multiobjective optimisation: A survey. *ACM Computing Surveys*, 52(2), 2019. doi: 10.1145/3300148 2, 3, 4
- [31] D. Liu, W. Cui, K. Jin, Y. Guo, and H. Qu. DeepTracker: Visualizing the training process of convolutional neural networks. *ACM Transactions on Intelligent Systems and Technology*, 10(1):1–25, 2018. doi: 10.1145/3200489 2
- [32] M. Liu, J. Shi, K. Cao, J. Zhu, and S. Liu. Analyzing the training processes of deep generative models. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):77–87, 2018. doi: 10.1109/TVCG.2017.2744938 2
- [33] S. Liu, J. Xiao, J. Liu, X. Wang, J. Wu, and J. Zhu. Visual diagnosis of tree boosting methods. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):163–173, 2017. doi: 10.1109/TVCG.2017.2744378 2
- [34] S. M. Lundberg, G. Erion, H. Chen, A. DeGrave, J. M. Prutkin, B. Nair, R. Katz, J. Himmelfarb, N. Bansal, and S. I. Lee. From local explanations to global understanding with explainable AI for trees. *Nature Machine Intelligence*, 2(1):56–67, 2020. doi: 10.1038/s42256-019-0138-9 2
- [35] Y. Ma, W. Chen, X. Ma, J. Xu, X. Huang, R. Maciejewski, and A. K. H. Tung. EasySVM: A visual analysis approach for open-box support vector machines. *Computational Visual Media*, 3(2):161–175, 2017. doi: 10.1007/s41095-017-0077-5 2
- [36] Y. Ma, A. Fan, J. He, A. R. Nelakurthi, and R. Maciejewski. A visual analytics framework for explaining and diagnosing transfer learning pro-

- cesses. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1385–1395, 2021. doi: [10.1109/TVCG.2020.3028888](https://doi.org/10.1109/TVCG.2020.3028888) 2
- [37] L. McInnes, J. Healy, N. Saul, and L. Grossberger. UMAP: Uniform manifold approximation and projection. *The Journal of Open Source Software*, 3(29):861, 2018. doi: [10.48550/arXiv.1802.03426](https://doi.org/10.48550/arXiv.1802.03426) 2
- [38] Y. Ming, H. Qu, and E. Bertini. RuleMatrix: Visualizing and understanding classifiers with rules. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):342–352, 2019. doi: [10.1109/TVCG.2018.2864812](https://doi.org/10.1109/TVCG.2018.2864812) 2
- [39] T. Muhlbacher, H. Piringer, S. Gratzl, M. Sedlmair, and M. Streit. Opening the black box: Strategies for increased user involvement in existing algorithm implementations. *IEEE Transactions on Visualization and Computer Graphics*, 20(12):1643–1652, 2014. doi: [10.1109/TVCG.2014.2346578](https://doi.org/10.1109/TVCG.2014.2346578) 2
- [40] M. A. Muñoz, M. Kirley, and S. K. Halgamuge. Exploratory landscape analysis of continuous space optimization problems using information content. *IEEE Transactions on Evolutionary Computation*, 19(1):74–87, 2015. doi: [10.1109/TEVC.2014.2302006](https://doi.org/10.1109/TEVC.2014.2302006) 8
- [41] D. Nagar, P. Ramu, and K. Deb. Visualization and analysis of Pareto-optimal fronts using interpretable self-organizing map (iSOM). *Swarm and Evolutionary Computation*, 76(August 2022):101202, 2022. doi: [10.1016/j.swevo.2022.101202](https://doi.org/10.1016/j.swevo.2022.101202) 2
- [42] J. L. J. Pereira, G. A. Oliver, M. B. Francisco, S. S. Cunha Jr, and G. F. Gomes. A review of multi-objective optimization: methods and algorithms in mechanical engineering problems. *Archives of Computational Methods in Engineering*, 29(4):2285–2308, 2022. doi: [10.1007/s11831-021-09663-x](https://doi.org/10.1007/s11831-021-09663-x) 1
- [43] N. Pezzotti, T. Höllt, J. Van Gemert, B. P. Lelieveldt, E. Eisemann, and A. Vilanova. DeepEyes: Progressive visual analytics for designing deep neural networks. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):98–108, 2018. doi: [10.1109/TVCG.2017.2744358](https://doi.org/10.1109/TVCG.2017.2744358) 2
- [44] A. Sarikaya, S. Member, and M. Gleicher. Scatterplots : Tasks, data, and designs. *IEEE Transactions on Visualization and Computer Graphics*, 24(1):402–412, 2018. doi: [10.1109/TVCG.2017.2744184](https://doi.org/10.1109/TVCG.2017.2744184) 9
- [45] H. Smedberg and S. Bandaru. Interactive knowledge discovery and knowledge visualization for decision support in multi-objective optimization. *European Journal of Operational Research*, 306(3):1311–1329, 2023. doi: [10.1016/j.ejor.2022.09.008](https://doi.org/10.1016/j.ejor.2022.09.008) 2
- [46] A. K. A. Talukder and K. Deb. PaletteStarViz: A visualization method for multi-criteria decision making from high-dimensional Pareto-optimal front. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 175–176, 2020. doi: [10.1145/3377929.3389939](https://doi.org/10.1145/3377929.3389939) 2
- [47] A. K. A. Talukder and K. Deb. PaletteViz: A visualization method for functional understanding of high-dimensional Pareto-optimal data-sets to aid multi-criteria decision making. *IEEE Computational Intelligence Magazine*, 15(2):36–48, 2020. doi: [10.1109/MCI.2020.2976184](https://doi.org/10.1109/MCI.2020.2976184) 2
- [48] Y. Tian, R. Cheng, X. Zhang, and Y. Jin. PlatEMO: A MATLAB platform for evolutionary multi-objective optimization. *IEEE Computational Intelligence Magazine*, 12(4):73–87, 2017. doi: [10.1109/MCI.2017.2742868](https://doi.org/10.1109/MCI.2017.2742868) 4, 6, 9
- [49] Y. Tian, L. Si, X. Zhang, R. Cheng, C. He, K. C. Tan, and Y. Jin. Evolutionary large-scale multi-objective optimization: A survey. *ACM Computing Surveys*, 54(8), 2021. doi: [10.1145/3470971](https://doi.org/10.1145/3470971) 2, 3, 5, 9
- [50] T. Tušar and B. Filipic. Visualization of Pareto front approximations in evolutionary multiobjective optimization: A critical review and the projection method. *IEEE Transactions on Evolutionary Computation*, 19(2):225–245, 2015. doi: [10.1109/TEVC.2014.2313407](https://doi.org/10.1109/TEVC.2014.2313407) 2
- [51] S. Van Den Elzen and J. J. Van Wijk. BaobabView: Interactive construction and analysis of decision trees. In *Proceedings of the IEEE Conference on Visual Analytics Science and Technology*, pp. 151–160, 2011. doi: [10.1109/VAST.2011.6102453](https://doi.org/10.1109/VAST.2011.6102453) 2
- [52] L. Van der Maaten and G. Hinton. Visualizing data using t-SNE. *Journal of Machine Learning Research*, 9(11), 2008. 2, 5
- [53] D. A. Van Veldhuizen, G. B. Lamont, et al. Evolutionary computation and convergence to a pareto front. In *Proceedings of the Genetic Programming Conference – Late-breaking Papers*, pp. 221–228, 1998. 5
- [54] M. J. Walter, D. J. Walker, and M. J. Craven. Visualising evolution history in multi- and many-objective optimisation. In *Parallel Problem Solving from Nature*, pp. 299–312, 2020. doi: [10.1007/978-3-030-58115-2_21](https://doi.org/10.1007/978-3-030-58115-2_21) 2, 3
- [55] M. J. Walter, D. J. Walker, and M. J. Craven. An explainable visualisation of the evolutionary search process. In *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, pp. 1794–1802, 2022. doi: [10.1145/3520304.3533984](https://doi.org/10.1145/3520304.3533984) 2, 3
- [56] M. J. Walter, D. J. Walker, and M. J. Craven. Visualizing population dynamics to examine algorithm performance. *IEEE Transactions on Evolutionary Computation*, 26(6):1501–1510, 2022. doi: [10.1109/TEVC.2022.3157143](https://doi.org/10.1109/TEVC.2022.3157143) 2, 3
- [57] J. Wang, L. Gou, H. W. Shen, and H. Yang. DQNViz: A visual analytics approach to understand deep Q-networks. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):288–298, 2019. doi: [10.1109/TVCG.2018.2864504](https://doi.org/10.1109/TVCG.2018.2864504) 2
- [58] J. Wang, S. Liu, and W. Zhang. Visual analytics for machine learning: A data perspective survey. *IEEE Transactions on Visualization and Computer Graphics*, pp. 1–20, 2024. doi: [10.1109/TVCG.2024.3357065](https://doi.org/10.1109/TVCG.2024.3357065) 2
- [59] Y. Wang, K. Feng, X. Chu, J. Zhang, C. W. Fu, M. Sedlmair, X. Yu, and B. Chen. A perception-driven approach to supervised dimensionality reduction for visualization. *IEEE Transactions on Visualization and Computer Graphics*, 24(5):1828–1840, 2018. doi: [10.1109/TVCG.2017.2701829](https://doi.org/10.1109/TVCG.2017.2701829) 2
- [60] Z. Wang, Y. Pei, and J. Li. A survey on search strategy of evolutionary multi-objective optimization algorithms. *Applied Sciences*, 13(7):4643, 2023. doi: [10.3390/app13074643](https://doi.org/10.3390/app13074643) 2, 6
- [61] D. Yadav, D. Nagar, P. Ramu, and K. Deb. Visualization-aided multi-criteria decision-making using interpretable self-organizing maps. *European Journal of Operational Research*, 309(3):1183–1200, 2023. doi: [10.1016/j.ejor.2023.01.062](https://doi.org/10.1016/j.ejor.2023.01.062) 2
- [62] J. Yuan, C. Chen, W. Yang, M. Liu, J. Xia, and S. Liu. A survey of visual analytics techniques for machine learning. *Computational Visual Media*, 7:3–36, 2021. doi: [10.1007/s41095-020-0191-7](https://doi.org/10.1007/s41095-020-0191-7) 2
- [63] J. Yuan, S. Xiang, J. Xia, L. Yu, and S. Liu. Evaluation of sampling methods for scatterplots. *IEEE Transactions on Visualization and Computer Graphics*, 27(2):1720–1730, 2021. doi: [10.1109/TVCG.2020.3030432](https://doi.org/10.1109/TVCG.2020.3030432) 9
- [64] Q. Zhang and H. Li. Moea/d: A multiobjective evolutionary algorithm based on decomposition. *IEEE Transactions on Evolutionary Computation*, 11(6):712–731, 2007. doi: [10.1109/TEVC.2007.892759](https://doi.org/10.1109/TEVC.2007.892759) 6
- [65] Y. Zhang and Q. Yang. A survey on multi-task learning. *IEEE Transactions on Knowledge and Data Engineering*, 34(12):5586–5609, 2021. doi: [10.1109/TKDE.2021.3070203](https://doi.org/10.1109/TKDE.2021.3070203) 1
- [66] X. Zhao, Y. Wu, D. L. Lee, and W. Cui. iForest: Interpreting random forests via visual analytics. *IEEE Transactions on Visualization and Computer Graphics*, 25(1):407–416, 2019. doi: [10.1109/TVCG.2018.2864475](https://doi.org/10.1109/TVCG.2018.2864475) 2
- [67] E. Zitzler, K. Deb, and L. Thiele. Comparison of multiobjective evolutionary algorithms: empirical results. *Evolutionary Computation*, 8(2):173–195, 2000. doi: [10.1162/106365600568202](https://doi.org/10.1162/106365600568202) 3