A Decision-Support Framework for Structuring and Reducing Large Multi-Objective Solution Sets via Clustering: An Application to Proton Therapy⋆

Nicole C. Cassimiro[a,*], Dylan Jones[b], Juliana C. Freitas[c] and Aurelio R. L. Oliveira[a]

[a]*Department of Applied Mathematics, Institute of Mathematics, Statistics and Scientific Computing, Universidade Estadual de Campinas, R. Sérgio Buarque de Holanda, 651, Campinas, 13083-859, São Paulo, Brazil*

[b]*School of Mathematics and Physics, University of Portsmouth, Portsmouth, PO1 3QJ, United Kingdom*

[c]*Department of Sciences, Universidade Estadual de Maringá, Goioerê, Paraná, Brazil*

ARTICLE INFO

Keywords:
Multi-objective optimization
Goal programming
Decision support
Clustering
OR for medicine

ABSTRACT

This paper proposes a four-stage decision-support framework for structuring and reducing large multi-objective solution sets into compact and interpretable collections of representative alternatives. The methodology combines: (Phase 1) systematic solution generation through extended goal programming and structured preference exploration; (Phase 2) robustness-aware enrichment and profiling under weight sensitivity analysis; (Phase 3) filtering and dominance-based reduction to remove infeasible or redundant alternatives; and (Phase 4) unsupervised clustering and representative selection to extract a small number of archetypal solutions. Rather than returning a single optimum, the framework constructs and analyzes ensembles of solutions and organizes them into meaningful trade-off profiles. The output is a concise decision subset that preserves the diversity of the original Pareto set while substantially reducing its size, thereby facilitating transparent and actionable decision-making. The approach is demonstrated on intensity-modulated proton therapy treatment planning for head-and-neck cancer, a naturally multi-objective and highly constrained problem that provides a demanding real-world testbed. Computational results show that large candidate ensembles can be consistently compressed into a small number of representative plans without loss of relevant trade-off information. The proposed methodology is shown to be scalable, and applicable to a broad class of multi-objective optimization problems.

1. Introduction

Large-scale multi-objective optimization problems frequently generate extensive sets of non-dominated or preference-sensitive solutions (Hartikainen, Miettinen and Klamroth, 2019; White, Witing, Wittekind, Volk and Strauch, 2025; Lotov, Bushenkov and Kamenev, 2017; Yadav, Nagar, Ramu and Deb, 2023). Although modern optimization techniques are effective at computing mathematically Pareto optimal alternatives, there is still room for research advances to support the organization, interpretation, and selection among these numerous competing solutions. In many practical settings, decision-makers are therefore confronted not with a single actionable answer, but with a large ensemble of trade-off solutions whose comparison and selection, in some fields, such as the medical planning domain, remain largely manual. This gap between optimization and decision-making creates a substantial cognitive and operational burden and motivates the development of systematic methodologies for structuring, reducing, and interpreting solution sets (Yadav et al., 2023).

Despite significant advances in multi-objective optimization, comparatively less attention has been devoted to post-optimization analysis and solution-set management. Classical scalarization techniques, such as weighted-sum formulations, require subjective parameter tuning and often yield solutions that depend strongly on ad hoc choices

This is a preprint of a manuscript submitted to the *European Journal of Operational Research*.

⋆ This document is the results of the research project funded by the Coordination for the Improvement of Higher Education Personnel (CAPES) and the National Council for Scientific and Technological Development – Brazil (CNPq).

*Corresponding author

nicole@ime.unicamp.br (N.C. Cassimiro); dylan.jones@port.ac.uk (D. Jones); jcfreitas@uem.br (J.C. Freitas); aurelio@ime.unicamp.br (A.R.L. Oliveira)

ORCID(s): 0000-0001-9263-1595 (N.C. Cassimiro); 0000-0002-9101-746X (D. Jones); 0000-0002-3147-4765 (J.C. Freitas); 0000-0002-6471-4710 (A.R.L. Oliveira)

(Marler and Arora, 2004). Even when more advanced goal-based or multi-criteria approaches are employed, the resulting models may produce many feasible or near-optimal alternatives whose trade-offs are difficult to assess in a transparent and reproducible manner. To bridge this gap, recent literature has seen important advances in explainability within multi-criteria decision analysis, seeking to make the evaluation and recommendation processes more transparent and understandable for the decision-maker (Corrente, Greco, Matarazzo and Słowiński, 2024; Corrente, Greco, Słowiński and Zappalà, 2026). Without incorporating such structured and explainable methodologies, optimization outputs are not always readily transformed into practical decision support (Greco, Słowiński and Wallenius, 2025).

Goal-based multi-objective optimization paradigms provide principled modeling frameworks for representing competing requirements and quantifying deviations from desired targets (Romero, 2004; Jones and Tamiz, 2010; dos Santos Rubem, Soares de Mello and Angulo Meza, 2017). In particular, Extended Goal Programming (EGP) stands out due to its specific strength: the parametric combination of the weighted and Chebyshev variants, which inherently represent efficiency and equity, respectively. Nevertheless, the processing and interpretation of the large ensembles of solutions that may result from such models, especially following a systematic preferential weight and parametric sensitivity analysis (Merchant, 2024), remain open challenges. Structured approaches for preference sensitivity analysis have been proposed in the goal and multiple-objective programming literature, notably through weight-sensitivity algorithms (Jones, 2011). However, these techniques are typically used to study local trade-offs and are rarely embedded within broader pipelines for large-scale solution generation, reduction, and decision support. Goal programming formulations have also been increasingly explored in complex decision support environments, such as healthcare and treatment planning contexts, where competing clinical criteria must be balanced in a structured and transparent manner (Malekpoor, Mishra and Kumar, 2022).

Moreover, advances in optimization and robustness analysis frequently increase, rather than decrease, the number of admissible solutions available to the decision-maker. Generating richer ensembles of candidate solutions improves the characterization of trade-offs but simultaneously intensifies the difficulty of comparison and selection. Methods capable of filtering dominated or redundant alternatives, identifying coherent families of solutions, and presenting compact sets of representatives are therefore essential for effective decision support (Lotov et al., 2017; White et al., 2025; An, Zhao and Zhu, 2025). Recent studies explicitly address pruning and reduction of large Pareto sets while preserving diversity and a decision-relevant structure (Petchrompo, Wannakrairot and Parlikad, 2022).

To address these challenges, this work proposes a novel operations research decision-support framework organized as a structured four-stage pipeline. The methodology explicitly separates solution generation from solution structuring: first producing a rich ensemble of preference-aware alternatives, and subsequently reducing this ensemble into a compact and interpretable decision subset. The four stages comprise: (i) systematic solution generation via extended goal programming and parametric preference exploration; (ii) robustness-aware enrichment and profiling of candidate solutions; (iii) filtering and dominance-based pruning to eliminate infeasible or redundant plans; and (iv) clustering-based structuring and representative selection to identify a small number of archetypal trade-off strategies.

To demonstrate the practical value of the approach, we apply the proposed framework to intensity-modulated proton therapy (IMPT) treatment planning, a domain that naturally exhibits highly conflicting objectives and complex trade-offs. Unlike conventional radiotherapy, which typically relies on photon (X-ray) beams, proton therapy is an advanced treatment modality that utilizes proton particles. A defining physical characteristic of protons is their ability to deposit the majority of their energy at a specific, controllable depth, known as the Bragg peak, and stop immediately thereafter, significantly sparing healthy tissues beyond the tumor. From a multi-criteria viewpoint, this planning process is formulated as a multi-objective optimization problem balancing target coverage, dose homogeneity, and organ-at-risk protection (Huiskes, Kong, Oud, Crama, Rasch, Breedveld, Heijmen and Astreinidou, 2024; Schubert and Teichert, 2025). While proton therapy offers highly conformal dose distributions, it is also highly sensitive to uncertainties and subject to strict dose-volume constraints and clinical reporting standards (Li, 2009; Sterpin, Widesott, Poels, Hoogeman, Korevaar, Lowe, Molinelli and Fracchiolla, 2024). These characteristics make IMPT a demanding and representative testbed for evaluating advanced decision-support methodologies (Mohamed, Li and Lee, 2022; Press and Mehta, 2024).

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

- A general decision-support framework for large multi-objective problems that integrates systematic solution generation, parametric preference exploration, robustness-aware enrichment, and clustering-based reduction into a unified four-phase pipeline that transforms extensive solution ensembles into compact and interpretable decision panels.

- A structured parametric exploration methodology that combines balance-parameter sweeps with hierarchical weight perturbations to generate diverse, preference-aware ensembles of candidate solutions and reveal the attainable compromise structure of the optimization model.
- A robustness-aware evaluation scheme based on quantitative performance and preference-stability metrics, enabling the characterization, comparison, and filtering of solutions beyond nominal objective values within the multi-criteria space.
- An unsupervised clustering and representative-selection methodology that identifies coherent families of trade-off solutions and extracts a small set of archetypal alternatives, effectively reducing large Pareto sets while preserving decision-relevant diversity.
- The development of a novel dose–volume EGP formulation tailored for IMPT inverse planning, integrating hierarchical targets, organ-specific dose limits, geometric rings, voxel-count dose–volume constraints, and voxel-wise Chebyshev homogeneity control within a unified mixed-integer framework. These constraint classes are structurally required to manage the ultra-steep spatial dose gradients of proton therapy and cannot be recovered by direct extension of photon-based EGP models. The formulation can be used both as a conventional single-plan optimizer and as a parametric solution generator underpinning the proposed framework.
- A comprehensive instantiation and evaluation of the framework using this challenging real-world IMPT testbed, demonstrating the effectiveness, scalability, and practical value of the methodology in complex decision-making environments.

2. Overview of the Proposed Decision-Support Framework

This section presents the proposed methodology. The objective is not to compute a single optimal solution, but to transform large multi-objective solution sets into compact and interpretable collections of representative alternatives that directly support decision-making.

Many real-world optimization problems produce extensive ensembles of Pareto-optimal or preference-sensitive solutions. Although such richness improves the characterization of trade-offs, it simultaneously increases the cognitive burden on decision-makers, who must compare numerous alternatives with similar performance. Consequently, effective decision support requires not only solution generation, but also systematic exploration, evaluation, reduction, and summarization of the solution space.

To address this need, we propose a unified framework that integrates optimization-based solution generation with robustness analysis and unsupervised machine learning techniques for solution-set reduction. This perspective aligns with recent decision-support research advocating interactive knowledge discovery and visualization pipelines that connect optimization outputs to actionable decisions (Smedberg and Bandaru, 2023).

The overall pipeline is summarized in Figure. 1 and is organized into four sequential phases.

- **Phase 1 – Structured preference exploration:** a goal-based multi-objective optimization model is repeatedly solved under systematic parameter and weight perturbations to map the compromise space and generate a diverse ensemble of candidate solutions;
- **Phase 2 – Solution enrichment and robustness profiling:** each solution is augmented with descriptive performance indicators, robustness signatures, and parametric coordinates, forming a structured solution space suitable for quantitative analysis;
- **Phase 3 – Clinical filtering and reduction:** clinically unacceptable or dominated solutions are removed through admissibility screening and dominance pruning, substantially reducing redundancy while preserving relevant trade-offs;
- **Phase 4 – Clustering and representative selection:** unsupervised machine learning techniques are applied to group similar solutions into coherent families, from which representative plans are selected to construct a compact and interpretable decision panel.

Although the proposed framework is general and applicable to a broad class of multi-objective optimization problems, in this work it is instantiated and evaluated on intensity-modulated proton therapy treatment planning, which serves as a demanding real-world testbed due to its inherently conflicting objectives and large candidate solution spaces.

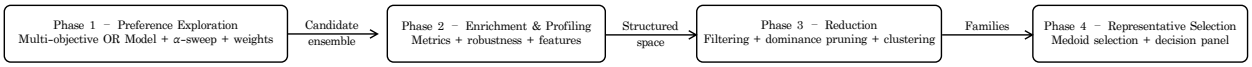

Figure 1: Four-stage information flow of the proposed decision-support framework.

3. Optimization Core: Goal-Based Model Formulation

The proposed decision-support framework requires an optimization core capable of generating high-quality solutions under competing objectives and being systematically perturbed to explore preference-sensitive trade-offs. To achieve this, we rely on EGP. While the modelling building blocks—deviation variables, hard feasibility limits, and parametric trade-off control—are domain-independent, existing EGP formulations in radiotherapy typically rely on simplified dose metrics or photon-based delivery assumptions (de Freitas, Cantane, Rocha and Dias, 2024; Schubert and Teichert, 2025). Conversely, IMPT-specific optimization models most often adopt weighted-sum or convex penalty formulations without an explicit goal-programming structure (Cao, Li, Chen, Wang, Si, Pei and Zhang, 2024; Kamal-Sayed, Ma, Tseung, Abdel-Rehim, Herman and Beltran, 2018). Although mixed-integer programming has recently been explored to enforce exact dose–volume constraints in IMPT (Zhang, Fan, Shan, Schild, Bues and Liu, 2017), such approaches have not yet been integrated into a comprehensive multi-objective framework capable of systematically balancing competing clinical goals.

The transition from photon radiotherapy to IMPT is not merely a substitution of the dose-influence matrix, but a fundamental shift in the mathematical structure of the optimization problem. Unlike conventional photons, IMPT utilizes pencil-beam scanning, requiring the exact intensity calibration of tens of thousands of individual microscopic proton spots. Furthermore, the sharp depth–dose fall-off characteristic of protons (the Bragg peak) enables highly conformal and homogeneous target dose distributions, but also produces ultra-steep spatial dose gradients that fundamentally alter how clinical objectives must be enforced in optimization. From a modelling perspective, photon-based EGP formulations can rely on smooth voxel penalties and approximate dose–volume surrogates, whereas IMPT planning requires explicit voxel-count dose–volume constraints, hierarchical target prescriptions, and ring-based spatial dose controls to manage proton-specific dose gradients. These constraint classes are therefore not optional refinements but structurally necessary components of an IMPT optimization model, and cannot be recovered by simply replacing the photon dose-influence matrix with its proton counterpart.

To bridge this gap, we present a unified dose–volume EGP formulation explicitly instantiated for IMPT. This model integrates hierarchical target prescriptions, voxel-wise Chebyshev homogeneity bounds, ring-based spatial fall-off constraints, and exact mixed-integer dose–volume endpoints directly into the spot-based inverse-planning environment. Beyond its role as the mathematical engine for our preference exploration framework, this formulation also serves as a rigorous standalone optimizer: for any fixed set of clinical priorities, it yields a single optimal IMPT plan consistent with the specified multi-criteria structure.

3.1. Notation and Clinical Problem Setting

To model the physical relationship between the scanning proton beam and the patient's anatomy, we define the fundamental sets, indices, and parameters that map the spatial dose distribution.

Typographical Conventions. To ensure clarity throughout the formulation, we adopt strict typographical rules: calligraphic uppercase letters (e.g., $\mathcal{V}, \mathcal{J}$) denote primary index sets; standard uppercase letters (e.g., V_T, D) denote subsets of voxels and matrices, respectively; lowercase subscripts (e.g., v, j, k) denote specific indices; and descriptive superscripts in Roman text (e.g., $d^{\text{pres}}, d^{\max}$) denote fixed clinical parameters or limits, distinguishing them from mathematical indices.

Sets and Indices. The problem is discretized into a set of delivery spots and a set of anatomical volume elements (voxels):

- $\mathcal{J} = \{1, \dots, N\}$: Set of deliverable proton pencil-beam spots, indexed by j.
- $\mathcal{V} = \{1, \dots, n\}$: Set of discrete 3D voxels representing the patient's body, indexed by v.

The anatomy is further partitioned into specific Regions of Interest (ROIs), representing different clinical priorities (Hodapp, 2012; International Commission on Radiation Units and Measurements, 2007; Breedveld and Heijmen,

2017). For any clinical structure s, let $V_s \subseteq \mathcal{V}$ be its specific subset of voxels, and $m_s = |V_s|$ be its volume (cardinality). The clinical structures are formally defined as:

- **CTV High**: The single primary tumor volume (labeled as T), represented by the voxel subset V_T, which must receive the full prescribed radiation dose to eradicate the disease.
- **Secondary Targets**: The set of adjacent areas at risk of microscopic spread, indexed by $k \in \mathcal{K}$, with corresponding voxel subsets $V_{T,k}$.
- **Organs-at-Risk (OARs)**: The set of serial organs at risk (e.g., spinal cord), indexed by $o \in \mathcal{O}_C$ with voxel subsets V_o^C, where exceeding a strict maximum dose limit causes irreversible damage.
- **Healthy Tissues**: The set of surrounding normal tissues and parallel organs, indexed by $o \in \mathcal{O}_H$ with voxel subsets V_o^H, where the goal is to minimize the overall average dose.
- **Rings**: The set of artificial structures, indexed by $r \in \mathcal{R}$ with voxel subsets V_r^R, constructed mathematically around the targets to force a sharp spatial dose drop-off in the optimization algorithm.

Accordingly, $m_T = |V_T|$, $m_{T,k} = |V_{T,k}|$, $m_o^C = |V_o^C|$, $m_o^H = |V_o^H|$, and $m_{R,r} = |V_r^R|$.

Decision Variables and Dose Operator. The main decision variable is the non-negative spot intensity vector $x = (x_j)_{j\in\mathcal{J}} \in \mathbb{R}_+^N$, which dictates the number of protons delivered to each individual spot. The physical dose deposition is governed by a dense dose-influence matrix $D \in \mathbb{R}^{n\times N}$, where D_{vj} is the absorbed dose (in Gray, Gy) delivered to voxel v by a unit intensity of spot j. The resulting nominal voxel-wise dose is exactly linear:

$$d_v(x) = \sum_{j\in\mathcal{J}} D_{vj} x_j \quad \forall v \in \mathcal{V}.$$

For any structure s, $D_s x$ denotes the dose vector restricted to the subset V_s, and the mean dose to that structure is $\bar{d}_s(x) = \frac{1}{m_s} \sum_{v\in V_s} d_v(x)$.

Clinical Prescriptions and Parameters. To guide the optimization, clinicians define specific numerical thresholds (parameters): prescriptions d_T^{pres} for the CTV High and $d_{T,k}^{\text{pres}}$ for secondary targets; maximum tolerated doses $d_o^{\max}$ for OARs; reference doses d_o^{ref} for OARs (used as mathematical normalization factors); and strict dose caps $d_{R,r}^{\max}$ for the rings.

3.2. Clinical Metrics and Goal-Based Requirements

A key modeling feature of our IMPT formulation is the exact mathematical encoding of selected dose–volume criteria within an EGP structure explicitly designed for parametric preference exploration. In large-scale IMPT, many multi-criteria formulations rely on continuous surrogates (e.g., voxel-wise penalties or mean-dose terms) to avoid the combinatorial burden of binary indicators (Kamal-Sayed et al., 2018). Here, we eschew these approximations to retain strict dose-volume constraints for the clinically decisive endpoints, while still enabling systematic control of efficiency–equity trade-offs through the EGP balance parameter and goal weights. The clinical logic is mathematically represented through three complementary mechanisms:

1. Hard requirements (strict limits). Non-negotiable safety limits, such as maximum doses to critical serial organs and ring drop-offs, are treated as absolute constraints (Li, 2009). These are modeled through strict voxel-wise caps to guarantee safety:

$$(D_o^C x)_v \le d_o^{\max} \quad \forall v \in V_o^C,\ \forall o \in \mathcal{O}_C; \qquad (D_r^R x)_v \le d_{R,r}^{\max} \quad \forall v \in V_r^R,\ \forall r \in \mathcal{R}.$$

2. Soft goals via deviations. Requirements that admit clinically meaningful trade-offs are formulated as goals with non-negative deviation variables. For a generic structure s and voxel $v \in V_s$, a prescribed target level b_s is linked to the delivered dose by $(D_s x)_v - p_{s,v} + n_{s,v} = b_s$, where $n_{s,v} \ge 0$ represents underdose (shortfall) and $p_{s,v} \ge 0$ represents overdose (excess). These deviations provide interpretable, linear measures of clinical violations and enable systematic weighting in the objective function.

3. Exact dose-volume requirements. In modern clinical protocols, treatments are strictly evaluated based on cumulative Dose-Volume Histograms (DVH). The standard evaluation metrics are defined as follows (Hodapp, 2012; Li, 2009):

- V_x: The percentage of a structure's volume receiving at least $x\%$ of the prescribed dose. For example, V_{95} measures target coverage, while V_{108} acts as a hot-spot control metric.
- D_y: The minimum dose received by the "hottest" $y\%$ of the structure's volume. D_{98} is a surrogate for the near-minimum dose (cold spot), and D_2 represents the near-maximu'm dose (hot spot).

Because enforcing exact V_x requirements fundamentally depends on counting the discrete number of voxels that satisfy prescribed dose conditions, these criteria are formulated using binary indicator variables (Zhang et al., 2017). To algebraically integrate this discrete counting logic into the continuous dose optimization, we employ a standard big-M formulation. This mixed-integer representation explicitly links the binary indicators to the voxel dose levels, embedding the non-convex DVH endpoints directly into the EGP model. This approach ultimately resolves the combinatorial challenge that conventional IMPT models bypass via penalty surrogates.

3.3. Extended Goal Programming Formulation

We now introduce the complete multi-objective, mixed-integer EGP formulation that constitutes the mathematical core of the proposed framework. This model integrates hierarchical target prescriptions, organ-specific dose limits, and exact dose-volume criteria within a single IMPT inverse-planning structure.

Decision Variables and EGP Parameters. To operationalize the clinical metrics defined in Section 3.2, the model employs four distinct classes of decision variables. Superscripts denote the associated anatomical structure, while subscripts index specific voxels or secondary targets:

- **Primary Physical Variables:** The continuous vector $x \in \mathbb{R}_+^N$ represents the deliverable proton spot intensities.
- **Continuous Deviations (Soft Goals):** Non-negative variables quantifying voxel-wise dose deviations (overdose p, underdose n) from prescribed or reference levels:
  - $p_v^T, n_v^T \geq 0$ for $v \in V_T$ (CTV High);
  - $p_{T,k,v}, n_{T,k,v} \geq 0$ for $v \in V_{T,k}$, $k \in \mathcal{K}$ (Secondary targets);
  - $p_{o,v}^H, n_{o,v}^H \geq 0$ for $v \in V_o^H$, $o \in \mathcal{O}_H$ (Healthy tissues).
- **Binary Indicators (Exact dose-volume):** Variables encoding the threshold-and-count logic for exact dose-volume criteria: $y_v^{\text{cov}} \in \{0,1\}$ and $y_v^{\text{hot}} \in \{0,1\}$ for the CTV High ($v \in V_T$); and $y_{k,v}^{\text{hot}} \in \{0,1\}$ for secondary targets ($v \in V_{T,k}$).
- **Auxiliary EGP Variables:** Two non-negative variables capture global clinical margins: δ_{cov} (the global coverage shortfall for CTV High) and $h_{\max}$ (the maximum homogeneity bound). Finally, the continuous scalar $\lambda \geq 0$ serves as the Chebyshev upper bound for all normalized goal components, mediating the balance between fairness (min-max) and efficiency (weighted sum) through the parameter α.

Clinical Priority Weights (Parameters). To steer the optimization toward clinically acceptable trade-offs, each goal component in the objective function is multiplied by a specific non-negative weight parameter, denoted by w. The nomenclature of these weights intuitively mirrors the clinical metrics they penalize: w_{pT} and w_{nT} correspond to overdose and underdose in the primary target; $w_{pT}^{(k)}$ and $w_{nT}^{(k)}$ apply to secondary targets; w_{cov} and w_{hom} penalize the global coverage shortfall and homogeneity margins; and w_{mean}^C and w_{mean}^H control the mean doses to critical and healthy organs, respectively. A comprehensive clinical interpretation of each weight is provided in Table 1.

Full EGP Model. The complete IMPT EGP formulation is presented below. Dose-volume requirements are encoded exactly using the binary indicator variables combined with standard big-M linking constraints, and the objective function systematically aggregates the weighted deviations.

$$\begin{aligned}
\min \quad & \alpha\,\lambda + (1-\alpha)\left[\sum_{v\in V_T}\left(\frac{w_{pT}}{m_T\,d_T^{\text{pres}}}\,p_v^T + \frac{w_{nT}}{m_T\,d_T^{\text{pres}}}\,n_v^T\right) + \frac{w_{\text{cov}}}{\rho_{\text{cov}}\,m_T}\,\delta_{\text{cov}} + \frac{w_{\text{hom}}}{d_T^{\text{pres}}}\,h_{\max}\right. \\
& + \sum_{k\in\mathcal{K}}\sum_{v\in V_{T,k}}\left(\frac{w_{pT}^{(k)}}{m_{T,k}\,d_{T,k}^{\text{pres}}}\,p_{T,k,v} + \frac{w_{nT}^{(k)}}{m_{T,k}\,d_{T,k}^{\text{pres}}}\,n_{T,k,v}\right) \\
& \left. + \sum_{o\in\mathcal{O}_C}\sum_{v\in V_o^C}\frac{w_{\text{mean}}^C}{m_o^C\,d_o^{\text{ref}}}\,(D_o^C x)_v + \sum_{o\in\mathcal{O}_H}\sum_{v\in V_o^H}\frac{w_{\text{mean}}^H}{m_o^H\,d_o^{\text{ref}}}\,(D_o^H x)_v\right] && \text{(1a)}
\end{aligned}$$

$$\text{s.t.} \quad \frac{w_{pT}}{m_T\,d_T^{\text{pres}}}\sum_{v\in V_T}p_v^T \le \lambda, \quad \frac{w_{nT}}{m_T\,d_T^{\text{pres}}}\sum_{v\in V_T}n_v^T \le \lambda, \tag{1b}$$

$$\frac{w_{\text{cov}}}{\rho_{\text{cov}}\,m_T}\,\delta_{\text{cov}} \le \lambda, \quad \frac{w_{\text{hom}}}{d_T^{\text{pres}}}\,h_{\max} \le \lambda, \tag{1c}$$

$$\frac{w_{pT}^{(k)}}{m_{T,k}\,d_{T,k}^{\text{pres}}}\sum_{v\in V_{T,k}}p_{T,k,v} \le \lambda, \quad \frac{w_{nT}^{(k)}}{m_{T,k}\,d_{T,k}^{\text{pres}}}\sum_{v\in V_{T,k}}n_{T,k,v} \le \lambda, \quad \forall\, k\in\mathcal{K}, \tag{1d}$$

$$\frac{w_{\text{mean}}^C}{m_o^C\,d_o^{\text{ref}}}\sum_{v\in V_o^C}(D_o^C x)_v \le \lambda, \quad \forall\, o\in\mathcal{O}_C, \tag{1e}$$

$$\frac{w_{\text{mean}}^H}{m_o^H\,d_o^{\text{ref}}}\sum_{v\in V_o^H}(D_o^H x)_v \le \lambda, \quad \forall\, o\in\mathcal{O}_H, \tag{1f}$$

$$(D_T x)_v - p_v^T + n_v^T = d_T^{\text{pres}}, \quad \forall\, v\in V_T, \tag{1g}$$

$$(D_{T,k} x)_v - p_{T,k,v} + n_{T,k,v} = d_{T,k}^{\text{pres}}, \quad \forall\, v\in V_{T,k},\ \forall\, k\in\mathcal{K}, \tag{1h}$$

$$(D_o^H x)_v - p_{o,v}^H + n_{o,v}^H = d_o^{\text{ref}}, \quad \forall\, o\in\mathcal{O}_H,\ \forall\, v\in V_o^H, \tag{1i}$$

$$(D_T x)_v \ge \theta_{\text{cov}}\,d_T^{\text{pres}} - M\,(1-y_v^{\text{cov}}), \quad \forall\, v\in V_T, \tag{1j}$$

$$\sum_{v\in V_T}y_v^{\text{cov}} + \delta_{\text{cov}} \ge \rho_{\text{cov}}\,m_T, \tag{1k}$$

$$(D_T x)_v \le \theta_{\text{hot}}\,d_T^{\text{pres}} + M\,y_v^{\text{hot}}, \quad \forall\, v\in V_T, \tag{1l}$$

$$\sum_{v\in V_T}y_v^{\text{hot}} \le \rho_{\text{hot}}\,m_T, \tag{1m}$$

$$(D_{T,k} x)_v \le \theta_{\text{hot}}\,d_{T,k}^{\text{pres}} + M\,y_{k,v}^{\text{hot}}, \quad \forall\, v\in V_{T,k},\ \forall\, k\in\mathcal{K}, \tag{1n}$$

$$\sum_{v\in V_{T,k}}y_{k,v}^{\text{hot}} \le \rho_{\text{hot}}\,m_{T,k}, \quad \forall\, k\in\mathcal{K}, \tag{1o}$$

$$(D_o^C x)_v \le d_o^{\max}, \quad \forall\, o\in\mathcal{O}_C,\ \forall\, v\in V_o^C, \tag{1p}$$

$$(D_r^R x)_v \le d_{R,r}^{\max}, \quad \forall\, r\in\mathcal{R},\ \forall\, v\in V_r^R, \tag{1q}$$

$$p_v^T \le h_{\max}, \quad n_v^T \le h_{\max}, \quad \forall\, v\in V_T, \tag{1r}$$

$$x_j \ge 0,\ \forall j\in\mathcal{J}; \quad p,n,h_{\max},\lambda,\delta_{\text{cov}} \ge 0; \quad y_v^{\text{cov}}, y_v^{\text{hot}}, y_{k,v}^{\text{hot}} \in \{0,1\}. \tag{1s}$$

The parameters θ_{cov} and ρ_{cov} define the target coverage requirement, ensuring that at least a fraction ρ_{cov} of the target volume receives at least θ_{cov} of the prescribed dose. Similarly, θ_{hot} and ρ_{hot} control hot-spot limitations by restricting the fraction of voxels that may exceed a specified dose level. In the clinical protocol used in this study, these parameters are set to

$$\theta_{\text{cov}} = 0.95, \quad \rho_{\text{cov}} = 0.98, \quad \theta_{\text{hot}} = 1.08, \quad \rho_{\text{hot}} = 0.02.$$

The formulation (1) encapsulates the clinical requirements into a unified MILP. The components of the model are structurally organized as follows:

- **Objective Function** (1a)**:** Combines a Chebyshev (min–max) component, $\alpha\lambda$, with a compensatory weighted-sum of normalized deviations. This structure directly regulates the degree of allowed trade-offs among the competing clinical goals.
- **Min–Max (Chebyshev) Bounds** (1b)**–**(1f)**:** Define the Chebyshev envelope. Each normalized goal component must be bounded by the continuous scalar λ. Consequently, λ represents the worst relative deviation across all clinical criteria, and its minimization promotes strictly balanced solutions.
- **Dose–Goal Tracking** (1g)**–**(1i)**:** Link the voxel-wise doses to the prescription or reference levels, defining the exact underdose (n) and overdose (p) continuous deviations for the primary target, secondary targets, and surrounding healthy tissues, respectively.
- **Exact Dose–Volume Criteria** (1j)**–**(1o)**:** Encode the non-convex DVH endpoints using a standard big-M mixed-integer formulation. Specifically, (1j) and (1k) strictly enforce target coverage by ensuring that at least a fraction ρ_{cov} of voxels in CTV High receives at least θ_{cov} of the prescribed dose (allowing for a global slack δ_{cov}).

  Constraints (1l)–(1o) limit the allowable fraction ρ_{hot} of voxels that may exceed θ_{hot} times the prescription dose (hot-spots) for both primary and secondary targets.
- **Hard Clinical Caps** (1p)**–**(1q)**:** Impose strict, uncompromisable voxel-wise dose limits for OARs and ring structures, guaranteeing absolute patient safety and enforcing sharp spatial dose fall-off.
- **Homogeneity Margin** (1r)**:** Defines the Chebyshev-type homogeneity control within CTV High by bounding both maximum local overdose and underdose deviations with the auxiliary scalar $h_{\max}$.
- **Variable Domains** (1s)**:** Establish the non-negativity of physical intensities and continuous deviations, as well as the binary nature of the exact threshold indicators.

The balance parameter $\alpha \in [0, 1]$ fundamentally controls the degree of compensation among the clinical goals. When $\alpha \approx 1$, the model prioritizes fairness by minimizing the worst normalized deviation (the Chebyshev regime), producing highly balanced plans where no single clinical requirement is excessively sacrificed. Conversely, when $\alpha \approx 0$, the objective approaches a pure weighted-sum aggregation, allowing for stronger compensatory trade-offs and favoring overall efficiency. This low-dimensional parametric control enables the systematic exploration of compromise solutions and underpins the α-sweep strategy adopted in our decision-support framework.

Finally, clinical priorities in the EGP model are expressed through a set of weights associated with each normalized deviation term in the compensatory weighted-sum component of the objective function. These weights encode the relative importance of coverage, homogeneity, inter-target hierarchy, and organ sparing, and reflect typical planning priorities in head-and-neck proton therapy. Beyond their clinical meaning, the weights play a central methodological role in this work: they constitute the preference parameters whose perturbation is used to analyze the stability and robustness of optimized plans. Table 1 summarizes the weights considered and their clinical interpretation.

Scope of the Optimization Core within the Framework. The formulation above specifies the optimization engine used to generate candidate solutions. From a modeling perspective, the primary focus of this framework is to isolate and resolve the multi-criteria decision-making challenges inherent in treatment planning. Consequently, our analysis of robustness centers strictly on *preference uncertainty*—that is, the stability of optimized plans under systematic perturbations of clinical priorities, objective weights, and trade-off parameters. To maintain this methodological focus on preference exploration, physical delivery uncertainties (e.g., patient setup and proton range errors) are not explicitly incorporated into the optimization, though scenario-based worst-case robustness could be integrated in the future by extending the dose-influence matrix. Furthermore, beam orientations and spot placements are assumed fixed *a priori*, isolating the intensity optimization and trade-off calibration from the combinatorial complexity of beam angle selection. Finally, all constraints and objective components are evaluated based on absorbed physical dose, leaving biological effect modeling (such as variable relative biological effectiveness or explicit complication probabilities) outside the current scope.

Table 1
Weights used in the EGP model and their clinical interpretation.

Component	Weight	Clinical interpretation
CTV High overdose	w_{pT}	Penalizes voxels in the high-priority target receiving dose above d_T^{pres}, reducing hot spots and supporting target uniformity.
CTV High underdose	w_{nT}	Penalizes voxels below d_T^{pres}, enforcing coverage and supporting tumor control.
Coverage shortfall (V_{95})	w_{cov}	Weights the global slack δ_{cov} associated with the coverage requirement (e.g., $V_{95} \geq 98\%$), discouraging plans that miss the desired covered volume.
Homogeneity margin	w_{hom}	Penalizes the Chebyshev-type homogeneity bound $h_{\max}$ that limits voxel-wise deviations around the prescription in CTV High, promoting a uniform therapeutic plateau.
Secondary target overdose	$w_{pT}^{(k)}$	Penalizes overdose in secondary targets to limit dose spillage from higher-dose regions and maintain strict spatial gradients across hierarchical target volumes.
Secondary target underdose	$w_{nT}^{(k)}$	Penalizes underdose in secondary targets to preserve their prescription levels $d_{T,k}^{\text{pres}}$ and maintain the intended target dose hierarchy.
OAR mean dose	w_{mean}^{C}	Minimizes mean dose to critical OARs (in addition to hard caps), providing further protection against severe toxicity. Normalized by d_o^{ref} for comparability across structures.
Healthy Tissue mean dose	w_{mean}^{H}	Minimizes mean dose to surrounding healthy tissues. Normalized by d_o^{ref}.

4. Proposed Decision-Support Framework

Building upon the mathematical core defined in Section 3, this section describes how the EGP engine is embedded into a structured decision-support pipeline. This framework systematically transforms large sets of preference-optimal solutions into compact, interpretable, and actionable decision alternatives.

4.1. Generation of the Candidate-Solution Ensemble

The EGP model depends on two classes of preference parameters: the balance parameter $\alpha \in [0, 1]$, which controls the trade-off between Chebyshev-type balancing and weighted aggregation; and a vector of weights w, which scales the normalized deviation terms.

Instead of tuning these parameters manually, we adopt a structured exploration of the preference space to generate a discrete yet representative ensemble of solutions. This strategy follows the principle that decision-makers should be presented with a spectrum of qualitatively different compromises rather than a single point estimate.

First, we perform an α-sweep over a uniform grid $\mathcal{A}$ defined by a discrete step size $\Delta\alpha$ (e.g., using $\Delta\alpha = 0.1$ to yield $\mathcal{A} = \{0, 0.1, 0.2, \ldots, 1.0\}$). This level of fidelity is chosen to strike a practical algorithmic balance: it provides sufficient granularity to capture non-linear shifts in the compromise spectrum while maintaining computational tractability for the large-scale MILP problem. For each $\alpha \in \mathcal{A}$, the EGP model is solved with a fixed baseline weight vector w, yielding:

$$x(\alpha) \in \arg\min \text{EGP}(\alpha, w).$$

This family of solutions approximates the attainable compromise spectrum of the model. A stable trade-off interval $\mathcal{A}^{\text{stab}} = [\alpha^-, \alpha^+] \cap \mathcal{A}$ is then identified as the subset of parameters for which solutions remain strictly feasible while exhibiting meaningful variation across competing objectives. Representative values $\{\alpha^-, \alpha^0, \alpha^+\}$, with $\alpha^0 = (\alpha^- + \alpha^+)/2$, are selected within this interval.

For each of these values, preference uncertainty is explored through systematic perturbations of the weight vector. Following the structured weight-sensitivity algorithm of Jones (2011), we generate a finite set $\mathcal{W}$ of perturbed weights using one- and two-dimensional scalings (considering a perturbation amplitude of $\rho = 0.5$, i.e., $\pm 50\%$ variations). Each combination (α, w) defines a unique optimization instance and yields a candidate solution $x(\alpha, w)$. In our IMPT

application, the resulting ensemble $\mathcal{P} = \{x(\alpha, w)\}$ constitutes a discrete, preference-structured approximation of the relevant deliverable treatment plans.

4.2. Filtering, Evaluation, and Robustness Ranking

To distill the generated ensemble $\mathcal{P}$ into a strictly relevant set of decision options, we apply a multi-tier post-optimization process. In many complex decision-making scenarios, the ultimate quality of a solution is judged by domain-specific indices that are computationally prohibitive to embed directly into the large-scale MILP optimization engine. Consequently, we separate the *optimization space* (the tractable EGP goals) from the *evaluation space* (the final evaluation metrics).

Let $\mathcal{K}$ be the set of relevant post-optimization evaluation domains for a given problem. We define a comprehensive normalized quality score $Q(x) \in [0, 1]$ for each candidate solution, computed as a weighted aggregation of these domain-specific indices:

$$Q(x) = \frac{\sum_{k\in\mathcal{K}} \beta_k q_k(x)}{\sum_{k\in\mathcal{K}} \beta_k},$$

where $q_k(x) \in [0, 1]$ represents the normalized performance sub-score for domain k, and β_k dictates its relative importance.

For the specific IMPT application addressed in this work, we define three core domains ($\mathcal{K} = \{\text{tumor}\ ,\ \text{OAR}\ ,\ \text{healthy}\}$). This yields the application-specific formulation:

$$Q(x) = \frac{\beta_T Q_{\text{tumor}}(x) + \beta_C Q_{\text{OAR}}(x) + \beta_H Q_{\text{healthy}}(x)}{\beta_T + \beta_C + \beta_H},$$

where Q_{tumor}, Q_{OAR}, and Q_{healthy} evaluate target coverage, critical structures maximum doses, and healthy tissue mean doses, respectively. In our baseline experiments, we adopt equal clinical priorities, setting $\beta_T = \beta_C = \beta_H = 1$.

While every plan in $\mathcal{P}$ is mathematically feasible and non-dominated within the specific configuration of the EGP model, it may exhibit structural inefficiencies or protocol violations when projected onto this newly defined performance space. Therefore, our post-optimization analysis proceeds as follows:

- **Acceptability Filtering:** We first impose minimal admissibility requirements corresponding to essential safety or target-coverage thresholds evaluated strictly *a posteriori*. Even though a solution is mathematically feasible in the MILP (due to the slack and deviation variables), it might be deemed unacceptable if it fails these explicit limits. Solutions violating these criteria are discarded, yielding a reduced acceptable subset $\mathcal{P}^{\text{acc}}$.
- **Dominance Pruning:** Next, strict dominance pruning is performed over the components of the subsequent evaluation vector $(q_k(x))_{k\in\mathcal{K}}$. Solutions from $\mathcal{P}^{\text{acc}}$ that are strictly dominated in this specific performance space are removed, producing a refined, non-dominated subset $\mathcal{P}^{\text{clean}}$.

Ultimately, each remaining solution in $\mathcal{P}^{\text{clean}}$ is characterized by its quality score $Q(x)$ and a robustness index $R(x) \in [0, 1]$. Here, $R(x)$ measures the solution's stability under the previously applied preference perturbations. Importantly, robustness is not used as a rigid cutoff filter, but rather as an evaluation metric to rank the surviving high-quality alternatives (which, in our specific application, correspond to the deliverable treatment plans).

To systematically rank these final alternatives, the quality and robustness scores are aggregated into the Parametric–Clinical Robustness Composite:

$$\text{PCRC}(x) = Q(x)^{\gamma} R(x)^{1-\gamma}, \qquad 0 < \gamma < 1.$$

By sorting the subset $\mathcal{P}^{\text{clean}}$ based on this metric, higher PCRC values explicitly identify the top-ranking solutions that achieve an optimal synergy between strong performance and stability against preference uncertainty.

4.3. Clustering and Structural Profiling of the Solution Space

Even after rigorous acceptability filtering and dominance pruning, the refined subset $\mathcal{P}^{\text{clean}}$ may still contain a large number of non-dominated alternatives. At this stage, the methodological objective shifts from optimization to structure discovery: we seek to identify coherent groups of solutions that represent qualitatively distinct compromise profiles.

To achieve this, each surviving solution is mapped into a normalized feature space that combines its multi-domain performance and stability:

$$z(x) = (q_1(x), q_2(x), \dots, q_{|\mathcal{K}|}(x), R(x)),$$

which jointly captures the competing evaluation metrics defined in Section 4.2 and the solution's robustness index.

Unsupervised clustering is then employed over this feature space as a general solution-set reduction mechanism. Specifically, Hierarchical Agglomerative Clustering (HAC) with Ward's linkage is adopted. This method is particularly well-suited for this task due to its ability to minimize within-cluster variance, producing highly compact, spherical groups, and its provision of an interpretable dendrogram representation (Ran, Xi, Lu, Wang and Lu, 2023; Ward, 1963; Murtagh and Contreras, 2012). For each identified cluster, a representative solution is selected using the centroid criterion (the point closest to the multidimensional center of the group). Together, these representatives form a compact, non-redundant decision subset that preserves the inherent diversity of the Pareto front. By analyzing the dominant features of each centroid, decision-makers can map these mathematically derived clusters into intuitive, domain-specific operational profiles.

In our IMPT case study, this multidimensional space translates to $z(x) = (Q_{\text{tumor}}, Q_{\text{OAR}}, Q_{\text{healthy}}, R)$. The resulting centroid representatives correspond directly to clinically distinct treatment strategies delivered to the physician for final selection. The integration of such unsupervised learning techniques into radiotherapy planning has gained significant traction recently, as clustering-based analyses have proven highly effective in grouping complex multi-criteria treatment plans to identify clinically consistent profiles and isolate structural anomalies (Huang, Yan, Song, Xu, Hu and Dai, 2023; Breedveld, Craft, van Haveren and Heijmen, 2017; Craft, Halabi, Shih and Bortfeld, 2012; Wu, Mohan, Morris and Lauve, 2019).

5. Computational Experiments and Results

We evaluate the proposed decision-support framework (detailed in Section 4 and summarized in Figure 1) using $N = 20$ head-and-neck IMPT clinical cases from the Radiotherapy Optimisation Test Set (TROTS) benchmark dataset (Breedveld and Heijmen, 2017). All optimization problems were formulated as mixed-integer linear programs and solved using Gurobi 10.0 on a standard workstation (Intel i7 CPU, 32 GB RAM) running macOS.

To provide transparency on the internal mechanics of the methodology and demonstrate how mathematically derived centroids translate into actionable clinical phenotypes, we first present a detailed, step-by-step walkthrough for a single representative instance (Section 5.1). Following this in-depth analysis, Section 5.2 reports cohort-level evidence to validate the framework's scalability, robustness, and consistency across the entire dataset.

5.1. Patient-Level Instantiation: Case Study P01

We instantiate the complete pipeline for patient `P01`. This instance is particularly suitable as a demonstration case due to the close geometric proximity between targets and multiple critical OARs, which induces pronounced and clinically challenging trade-offs. Such characteristics naturally generate a large and high-dimensional solution space, making the case a demanding yet informative benchmark instance.

Following the hierarchical target definition established in Section 3, the specific EGP formulation for this case includes a high-dose clinical target volume (CTV High) prescribed to 64.68 Gy, alongside secondary target volumes (CTV Intermediate and CTV Low Shrunk) prescribed to 52.92 Gy. These volumes are adjacent to critical serial structures, namely the spinal cord and the brainstem, for which a strict maximum-dose limit of 20 Gy is enforced.

In addition to these critical organs, the model incorporates multiple dose-sensitive healthy structures relevant for post-treatment quality of life: the larynx, the oral cavity, and the bilateral submandibular and parotid glands. For these organs, mean-dose objectives are dynamically utilized to promote maximum dose sparing while maintaining strict target coverage. Finally, geometric ring structures surrounding the target volumes (e.g., CTV High Ring at 0–10 mm) are included as auxiliary structures. These rings enforce a sharp spatial dose fall-off, limiting excessive dose spillage into nearby healthy regions.

The experimental results for this case are presented below in direct correspondence with the pipeline stages previously established in Section 4.

5.1.1. Phase 1: Compromise Mapping via the α-Sweep

To map the attainable compromise space for patient `P01`, the EGP optimization core was solved across the prescribed grid $\mathcal{A}$. This one-dimensional α-sweep systematically generates a structured baseline family of solutions,

spanning the analytical spectrum between strict target conformity ($\alpha \to 1$, Chebyshev minimax regime) and aggressive organ sparing ($\alpha \to 0$, pure weighted-sum regime). Figure 2 summarizes the resulting trade-off trajectories, reporting target coverage, critical OAR maximums, healthy tissue mean doses, and computational runtimes as functions of α.

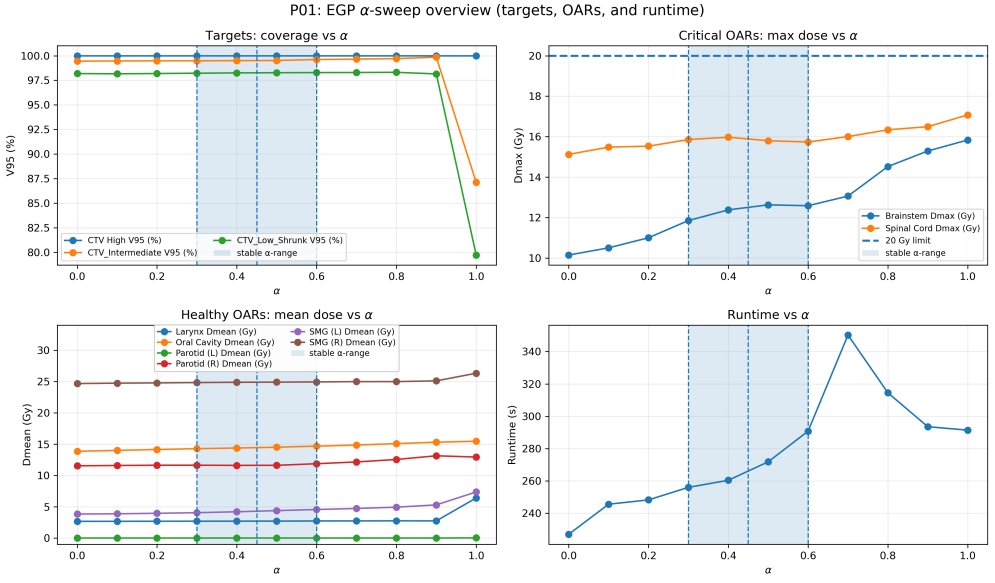

Figure 2: Patient-level α-sweep analysis for case P01. Top left: target coverage (V_{95}) for the high-, intermediate-, and low-dose target volumes as a function of α. Top right: maximum dose to critical OARs (spinal cord and brainstem), with the clinical limit indicated by the dashed line. Bottom left: mean dose to healthy OARs. Bottom right: computational runtime per optimization. The shaded region highlights the stable trade-off interval ($\mathcal{A}^{\text{stab}}$) identified for subsequent preference exploration.

Two complementary observations arise from this mapping. First, the strictly constrained metrics—target coverage and OAR protection—remain remarkably stable across a broad intermediate range of α. In this regime, the target volumes maintain near-complete coverage, and maximum doses to the spinal cord and brainstem remain strictly below their absolute clinical limits. This indicates that the core feasibility requirements of the EGP model are highly robust to moderate variations in the scalarization balance.

Second, the active multi-objective trade-offs are predominantly observed in the unconstrained healthy tissues. In these structures, dose redistribution reflects the expected mathematical competition between spatial conformity and general tissue sparing. Importantly, these objective values evolve smoothly with α, without abrupt regime shifts or extreme discontinuities. This confirms that the EGP formulation yields a continuous and highly stable compromise landscape, rather than isolated or hypersensitive solutions.

Together, these empirical findings identify the patient-specific *stable trade-off interval*, denoted as $\mathcal{A}^{\text{stab}}$ (highlighted in Figure 2). From a methodological standpoint, establishing this interval is a critical prerequisite: it rigorously bounds the relevant decision space, ensuring that the computationally intensive weight perturbations in Phase 2 are exclusively focused on clinically viable and mathematically stable regions of the Pareto front.

5.1.2. Phase 2: Preference Sensitivity and Robustness Profiling

Within the established stable interval $\mathcal{A}^{\text{stab}}$, systematic one- and two-dimensional weight perturbations were applied to the EGP baseline configuration. For each perturbed instance, the optimization core was re-solved, and the resulting solutions were evaluated using the parametric stability index R and the composite metric PCRC. Figure 3 visualizes the directional impact of these weight variations on the objective values, robustness indicators, and computational runtimes.

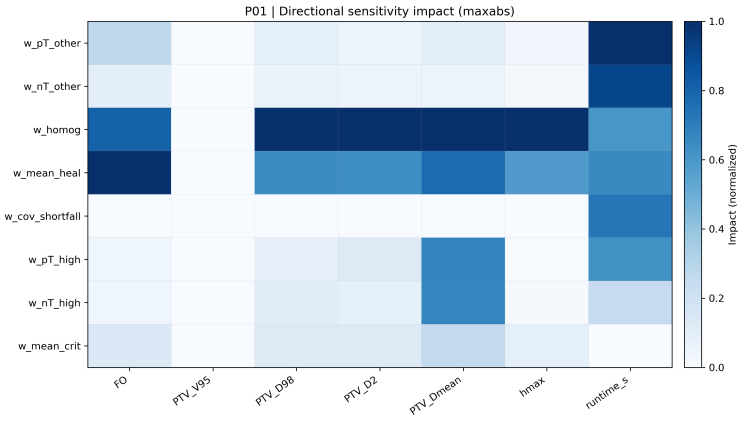

Figure 3: Directional sensitivity analysis for patient P01. Column-normalized impacts highlight which preference weights most strongly influence plan quality, robustness, and runtime within the stable α interval.

The empirical sensitivity structure reveals a highly non-uniform influence of the objective weights. Perturbations associated with target homogeneity and healthy-tissue mean doses act as the primary drivers of dose redistribution, producing the largest variances in both the final solution vectors and the solver runtimes. In contrast, increasing the penalty weights for target underdose yields diminishing returns; once initial coverage thresholds are satisfied, these components become largely insensitive to further penalization.

This behavior highlights a clear mathematical decoupling between *feasibility* and *Pareto-front refinement*. The coverage requirements effectively act as hard admissibility bounds, whereas the homogeneity and normal-tissue sparing objectives govern the continuous geometric shaping of the dose distribution. Crucially, despite these distinct directional sensitivities, the overall robustness scores remain strictly controlled. The candidate solutions do not exhibit abrupt degradation under moderate preference shifts.

Furthermore, an analysis of the two-dimensional weight variations confirmed that interaction effects between competing clinical priorities do not introduce hidden instability mechanisms. The resulting sensitivity patterns perfectly mirror the one-dimensional analysis. Consequently, this phase successfully yields a densely populated, preference-stable solution space, rigorously primed for the post-optimization reduction and clustering stages.

5.1.3. Phase 3: Acceptability Filtering and Dominance Pruning

The systematic preference sweeps generated an initial ensemble $\mathcal{P}$ of 210 candidate solutions for patient P01. To distill this dataset into a clinically relevant selection, we applied a two-stage post-optimization reduction pipeline.

First, the acceptability filter discarded 18 solutions that violated strict *a posteriori* limits for critical OARs or minimum target coverage, resulting in an intermediate subset $\mathcal{P}^{\text{acc}}$ of 192 plans. While these plans are all clinically deliverable, they contain significant redundancy and varying levels of dosimetric trade-offs.

Next, a multi-criteria dominance pruning was executed over the clinical performance space defined by (Q_{tumor} , Q_{OAR} , Q_{healthy}). This step revealed that most of the 192 clinically acceptable plans were Pareto-dominated by other candidates and therefore redundant from a multi-objective perspective. The dominance filtering reduced the set to a refined subset $\mathcal{P}^{clean}$ consisting of 14 non-dominated plans.

This reduction of over 90% of the acceptable search space highlights the efficacy of the framework: most weight combinations in the optimization engine converge to similar or inferior dosimetric outcomes. By isolating these 14 optimal trade-offs, the framework focuses only on the truly distinct clinical strategies, which are subsequently organized into 4 archetypal representatives through the final clustering stage.

5.1.4. Phase 4: Clustering, Clinical Profiling, and Representative Selection

To uncover the underlying structural typology of the 14 non-dominated plans in $\mathcal{P}^{\text{clean}}$, we applied HAC over the normalized clinical–robustness feature space $z(x) = (Q_{\text{tumor}}\ ,\ Q_{\text{OAR}}\ ,\ Q_{\text{healthy}}\ ,\ R)$. As illustrated by the dataset's dendrogram and the corresponding elbow criterion (Figure 4), this unsupervised analysis identified $k = 4$ optimal families. This confirms a clear hierarchical taxonomy of trade-off strategies rather than a diffuse continuum.

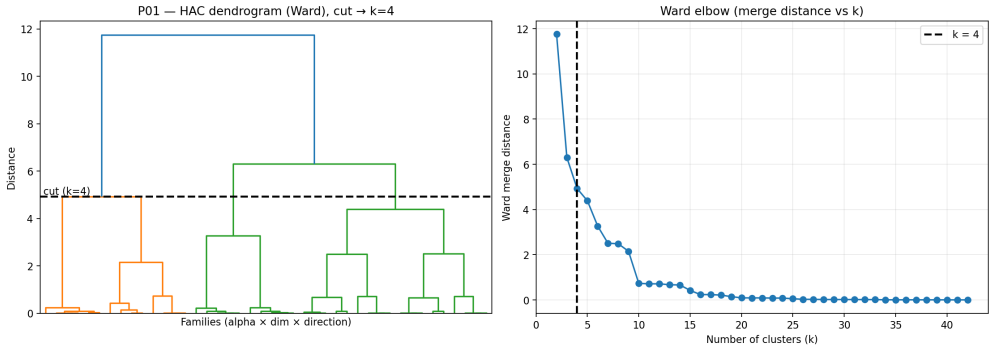

Figure 4: Hierarchical clustering structure for patient P01. The dendrogram illustrates the multidimensional separation of the 14 non-dominated plans, while the elbow method analysis (inset/side panel) justifies the selection of $k = 4$ archetypal families.

To visualize this high-dimensional separation, Figure 5 projects the candidate solutions into a two-dimensional UMAP embedding (employed exclusively as an interpretive visualization tool, not for the clustering itself). To provide a meaningful external baseline, we also project a standard treatment plan obtained from the classical nominal weighted-sum scalarization commonly adopted in proton therapy (Cao, Romeijn, Shepard and Craft, 2012).

The resulting macroscopic structure is highly informative: the admissible plans organize into tightly clustered, distinct regions. The classical nominal solution (marked as a black X) appears as merely a single, arbitrary point within this structured landscape. This illustrates a fundamental limitation of the conventional single-plan approach: it samples only one isolated location of a much richer trade-off space, whereas the proposed methodology systematically exposes its full, mathematically viable organization.

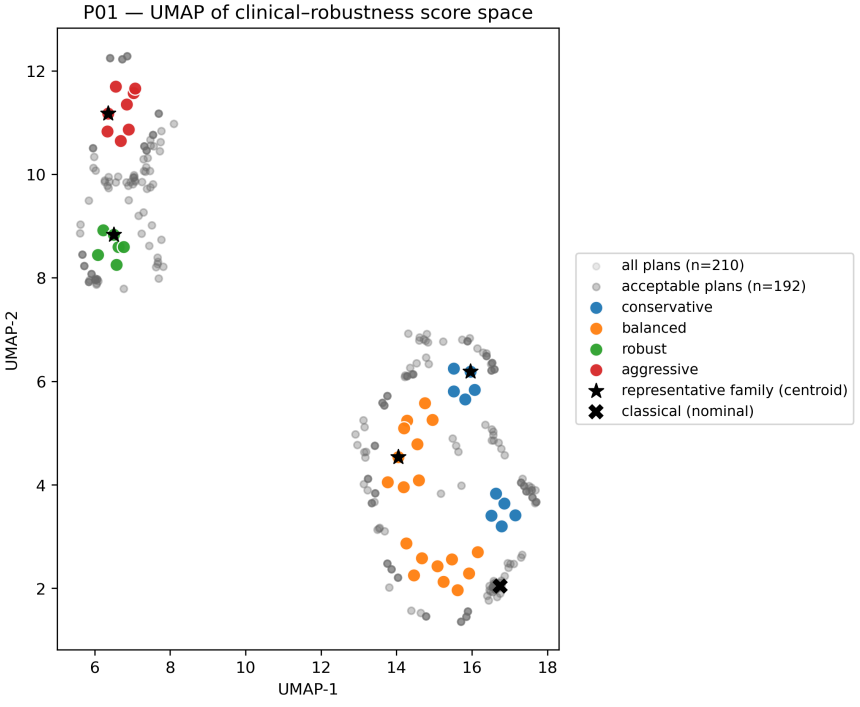

Figure 5: UMAP visualization of the clinical–robustness score space for patient P01, showing all candidate plans, the acceptable subset, the HAC families ($k = 4$), the selected representatives (black stars), and the classical nominal weighted-sum solution (black X). The clustering is performed in the original standardized multi-dimensional space.

Following the methodology established in Section 4.3, a single representative plan (medoid) is extracted from each cluster by selecting the solution closest to the mathematical centroid. This choice guarantees the selection of stable, archetypal behavior, avoiding outliers that might score highly on a single metric but lack overall robustness.

These $k = 4$ representatives form a compact, non-redundant decision panel that compresses the entire admissible space. Clinically, these mathematically derived centroids translate directly into distinct, actionable phenotypes (e.g., aggressive tumor-focused, strictly conservative, balanced, and robustness-oriented profiles).

Phase 4 thus completes the transition from a large, high-dimensional ensemble to a concise set of alternatives. However, structural separation in the mathematical feature space does not, by itself, establish clinical superiority. Since decision support ultimately depends on the physical dose distribution, we immediately complement this score-space analysis with a direct dosimetric evaluation. The subsequent section compares these four clustered representatives against the classical nominal plan using standard Dose–Volume Histograms (DVH), verifying the tangible clinical value of the proposed decision subset.

5.1.5. Dose–Volume Histogram Analysis and Clinical Interpretation

To evaluate the practical clinical value of the generated decision subset, we project the selected representative plans back into the physical dose domain. We examine whether the multi-objective structured panel translates into tangible improvements over the classical single-solution planning strategy. The analysis proceeds in three steps: (i) a visual DVH comparison, (ii) a quantitative assessment of relative improvements, and (iii) the extraction of clinical fingerprints for each representative family.

Nominal versus Robust Representative Plan. We first compare the classical nominal plan obtained via standard weighted-sum scalarization (Cao et al., 2012) against the medoid of the *Robust* family extracted by our framework. Figure 6 reports the DVHs for both plans. In a DVH, the ideal target curve resembles a steep step function at the prescribed dose level, while optimal OAR curves shift as far to the left (lower dose) as possible.

As shown, target coverage for the high-, intermediate-, and low-dose CTVs remains essentially identical. Crucially, both solutions fully satisfy the primary clinical acceptability criteria mathematically enforced by the EGP formulation: delivering at least 95% of the prescribed dose to 98% of the target volume ($D_{98} \geq 95\%$), while strictly limiting the maximum dose hotspot to 107% ($D_2 \leq 107\%$).

With these core tumor-control requirements perfectly secured, pronounced systematic differences emerge in the healthy tissues. For non-clinical readers, these leftward shifts in the DVH represent major sparing gains. For example, reading the curves at the 30% volume threshold for the right parotid gland (Parotid R), the robust plan delivers approximately 15 Gy, whereas the nominal plan delivers nearly 25 Gy to the same volume fraction. Similarly, for the larynx, the robust plan strictly limits the dose to 10 Gy for 10% of the volume, while the nominal plan allows up to 20 Gy.

These improvements emerge naturally from the structured preference exploration, demonstrating that the framework identifies solutions that preserve strict tumor coverage while drastically improving normal-tissue sparing and stability.

Quantitative Comparison with the Classical Nominal Plan. To quantify these gains, Table 2 reports the direction-adjusted relative percentage improvement of all four family representatives against the classical nominal baseline. Positive values strictly indicate clinically favorable changes (higher target coverage or lower OAR dose).

Across multiple critical and healthy structures, the framework's representatives achieve systematic, double-digit percentage gains. For instance, the mean dose reductions in the larynx (approx. 60%) and critical maximum dose reductions in the spinal cord (over 23%) represent massive clinical benefits. This proves that the classical scalarization converges to a highly suboptimal region of the attainable trade-off surface, whereas the proposed methodology reliably extracts quantitatively superior alternatives.

Clinical Profiles Across Representative Families. While the representatives uniformly outperform the nominal baseline, absolute dose variations (in Gy) among Pareto-optimal plans can appear numerically small, despite being highly significant to a clinician. To rigorously characterize the internal trade-off structure among the representatives themselves, Table 3 reports direction-adjusted min–max normalized scores (0–1). Here, 1 corresponds to the best-performing family for a given metric across the generated panel.

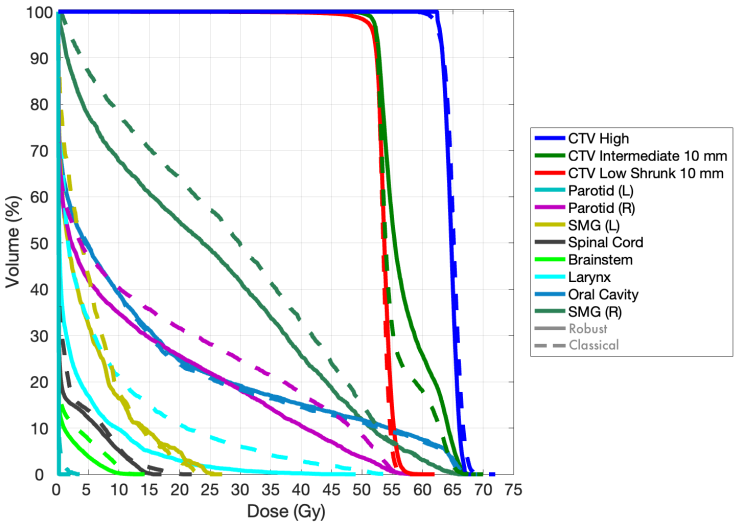

Figure 6: Dose–volume histograms for patient P01, comparing the classical nominal plan (dashed lines) and the representative robust plan (solid lines). While target coverage is virtually identical, the robust plan significantly shifts OAR curves to the left, demonstrating improved intermediate-dose sparing (e.g., Parotid glands and Larynx).

Table 2
P01: Direction-adjusted improvement (%) of family representatives relative to the Nominal (Cao et al. (2012) weighted-sum) baseline. Positive values indicate clinically favorable changes.

Family	**Tumor** (%)		**OARs** (%)		**Healthy Tissues** (%)		
	D_{98}	D_2	SC $D_{\max}$	BS $D_{\max}$	Larynx D_{mean}	SMG (R) D_{mean}	Parotid (R) D_{mean}
Aggressive	+0.92	+1.27	+23.93	+16.03	+59.36	+14.06	+20.90
Balanced	+0.74	+1.20	+23.35	+20.93	+59.95	+14.41	+22.49
Robust	+0.92	+1.27	+23.93	+16.03	+59.36	+14.06	+20.90
Conservative	+0.74	+1.20	+23.35	+20.93	+59.95	+14.41	+22.49

Table 3
P01: Min–max normalized clinical fingerprint (0–1). Scores are computed strictly across the four family representatives. A value of 1 indicates the best phenotype for that specific metric; values isolate the relative trade-off strategy rather than absolute feasibility.

Family	**Tumor**		**OARs**		**Healthy Tissues**			**Robustness**	
	D_{98}	D_2	SC $D_{\max}$	BS $D_{\max}$	Larynx D_{mean}	SMG (R) D_{mean}	Parotid (R) D_{mean}	R	PCRC
Aggressive	1.00	1.00	1.00	0.00	0.00	0.00	0.00	0.00	0.00
Balanced	0.00	0.00	0.00	1.00	1.00	1.00	1.00	0.14	0.47
Robust	1.00	1.00	1.00	0.00	0.00	0.00	0.00	1.00	0.57
Conservative	0.00	0.00	0.00	1.00	1.00	1.00	1.00	1.00	1.00

This normalization acts as a compact clinical *fingerprint*, revealing the distinct therapeutic phenotype of each cluster. The Aggressive family maximizes target conformity but sacrifices normal tissue; the Conservative family provides maximum healthy organ protection; the Balanced family sits in the geometric middle; and the Robust family matches the dosimetric quality of the Aggressive plan while drastically improving stability against preference perturbations (scoring 1.00 in the R index).

These structured and interpretable profiles confirm that the feasible region is not redundant, but rather mathematically organized into a small number of clinically meaningful strategies. Consequently, the representative selection stage acts not merely as dimensionality reduction, but as an effective decision-support mechanism that exposes genuinely distinct therapeutic options to the physician.

5.2. Cohort-Level Validation: Scalability, Robustness, and Set Reduction

Having demonstrated the internal mechanics and clinical interpretability of the decision-support framework for a single complex instance (P01), we now expand our analysis to the full cohort. The objective of this section is to validate the scalability, consistency, and structural robustness of the proposed methodology across all $N = 20$ head-and-neck IMPT clinical cases from the TROTS dataset.

As established, each instance represents a highly complex multi-objective EGP formulation comprising hierarchical targets and multiple critical structures.

For a typical patient, the baseline Phase 1 α-sweep (11 scalarization solves) required an average of 50.8 minutes in total, corresponding to 4.62 minutes per solve (median 4.53; range 3.78–5.84). The Phase 2 preference sensitivity stage naturally constitutes the dominant computational cost of the pipeline. The systematic one- and two-dimensional weight perturbations required a standardized grid of 199 additional optimization runs per patient. This generative phase required an average total computation time of 46.30 hours per patient (median 45.59; range 19.03–77.62).

In strict contrast, the subsequent post-optimization stages, acceptability filtering, dominance pruning, robustness scoring, clustering, and representative selection, rely exclusively on deterministic linear-algebraic and data-analytic operations. Consequently, Phases 3 and 4 account for a negligible fraction of the total runtime (less than 4 seconds per patient).

This profound structural separation between computationally intensive solution generation (Phases 1 and 2) and lightweight structuring (Phases 3 and 4) is a deliberate design choice. It guarantees that once the offline mathematical generation of the Pareto-front approximation is complete, the actual clinical decision-support process (filtering and grouping) can be performed almost instantaneously.

5.2.1. Phase 1: Structured Preference Exploration via the α-Sweep

To evaluate the overarching macroscopic behavior of the baseline EGP formulation, we executed a structured sweep of the balance parameter $\alpha \in [0, 1]$ across all $N = 20$ patients. For each instance, the model was solved over the discrete grid $\mathcal{A} = \{0, 0.1, \ldots, 1.0\}$, and the resulting plans were mapped into the normalized quality indices Q_{tumor}, Q_{OAR}, and Q_{healthy}.

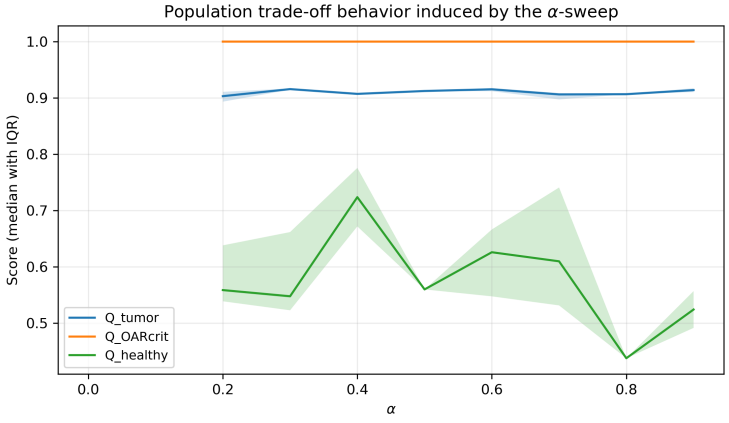

Figure 7: Cohort-level trade-off profiles along the α-sweep (Phase 1 baseline plans). Solid lines denote the median across all 20 patients, and shaded regions represent the interquartile range (25th–75th percentiles). The notably flat trajectories demonstrate the parametric stability of the EGP formulation.

The empirical α-sweep reveals a highly *stable*, rather than geometrically volatile, trade-off landscape. As shown in Figure 7, three consistent patterns emerge across the cohort:

First, the compromise surface is fundamentally continuous and highly insensitive to the α parameter, lacking any abrupt regime shifts. Second, Q_{tumor} and Q_{OAR} remain virtually invariant across the entire grid. This indicates that the solution space is organized into a small number of structurally distinct trade-off regimes rather than forming a diffuse continuum of gradual variations. Third, Q_{healthy} exhibits the only noticeable (yet entirely smooth) dependency on α, confirming that preference variations primarily govern the secondary redistribution of the low-dose bath in non-critical healthy tissues.

Collectively, these observations delineate a broad *stability plateau*. Across the cohort, clinically acceptable and mathematically sound plans can be reliably extracted with limited sensitivity to the exact choice of α. This phenomenon confirms the existence of a robust, patient-specific stable interval $\mathcal{A}^{\text{stab}}$, providing the foundation for the multi-dimensional preference perturbations executed in Phase 2.

5.2.2. Phase 2: Preference Sensitivity and Robustness Profiling

Anchored within the stability plateau established in Phase 1, Phase 2 systematically executed one- and two-dimensional weight perturbations across the full $N = 20$ cohort. This phase transitions the analysis from global objective balancing (the α-sweep) to localized sensitivity profiling, rigorously testing how the underlying EGP solutions respond to clinical preference uncertainty.

For the resulting ensemble of clinically acceptable plans, parametric stability was quantified using the robustness index R and the composite score PCRC. As depicted in Figure 8, the cohort-level distribution of R exhibits a pronounced lower tail. This is a critical finding: it reveals the latent presence of highly preference-sensitive (brittle) solutions that would be indistinguishable from robust plans under standard nominal evaluation. In contrast, the PCRC distribution is markedly tighter and right-shifted. By explicitly penalizing both dosimetric inadequacy and mathematical instability, PCRC effectively filters out these brittle outliers, isolating the robust core of the Pareto approximation.

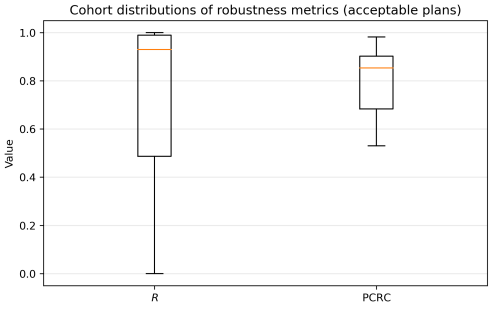

Figure 8: Cohort-level distributions of robustness metrics for clinically acceptable plans. The composite PCRC metric exhibits a higher concentration and stricter central tendency than the isolated R index, effectively filtering out brittle solutions.

The impact of perturbation dimensionality is further dissected in Figure 9. While isolated one-dimensional (1D) perturbations generally yield high and concentrated stability scores, exposing the models to simultaneous two-dimensional (2D) weight variations naturally increases dosimetric variability and slightly reduces the central tendency.

Crucially, however, the robustness metrics demonstrate *graceful degradation* rather than a catastrophic collapse under multi-parameter uncertainty. The absence of severe performance cliffs confirms that the EGP formulation's feasible region possesses inherent structural integrity.

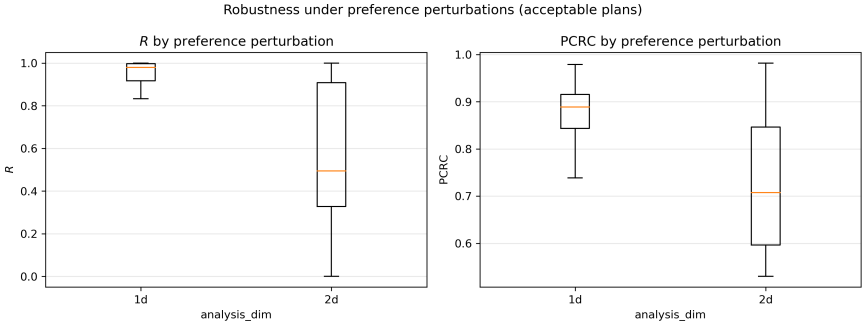

Figure 9: Robustness score distributions under preference perturbations across the cohort, grouped by perturbation complexity (1D vs. 2D variations). The system demonstrates graceful degradation under multi-parameter uncertainty.

Ultimately, Phase 2 confirms that the generated candidate ensembles are not merely mathematically diverse, but possess a controlled, rigorously quantifiable sensitivity to preference shifts. This robustness-aware profiling provides the definitive data foundation required for the subsequent dimensionality reduction and clustering stages.

5.2.3. Phase 3: Post-Optimization Reduction of the Candidate Ensemble

The structured preference exploration in Phases 1 and 2 generates a high-density ensemble, yielding a median of 210 candidate plans per patient (range: 209–210). Phase 3 applies the post-optimization reduction pipeline to compress this extensive dataset into an actionable panel.

The quantitative impact of this progressive compression is summarized in Figure 10. First, the clinical admissibility filter removed solutions that violated strict *a posteriori* OAR limits or tumor coverage threshold. This stage proved highly selective, refining the cohort-level pool to a median of 41.5 clinically acceptable plans per patient. This 80% reduction rate underscores the importance of the filtering stage in discarding mathematically valid but clinically sub-optimal trade-offs.

Next, Pareto-dominance pruning was executed across the four-dimensional performance space (Q_{tumor}, Q_{OAR}, Q_{healthy}, and R). This operation effectively eliminated remaining redundancies, converging to a median of 4.0 non-dominated solutions.

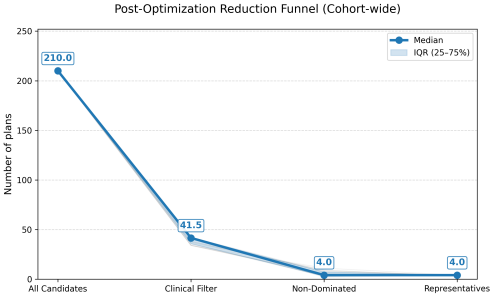

Figure 10: Phase 3 reduction pipeline: progressive compression of candidate plans across the $N = 20$ cohort. Data labels indicate the median number of plans at each stage (Initial: 210.0 → Clinical Filter: 41.5 → Non-Dominated: 4.0 → Representatives: 4.0). The pipeline isolates the most robust and clinically superior options from the initial generative ensemble.

Finally, the clustering-driven selection confirmed this structural concentration by identifying 4.0 archetypal representatives. As demonstrated, the framework successfully reduces the decision space by nearly two orders of magnitude from 210 computationally generated solutions down to exactly 4 distinct, robust clinical options. This concentration highlights the efficacy of the pipeline in isolating the high-quality core of the Pareto approximation while maintaining clinical feasibility.

5.2.4. Phase 4 — Clinical Profiling, Clustering, and Archetype Selection

Phase 4 moves beyond data reduction to the structural interpretation of the remaining solution space. The primary finding is that the non-dominated space is organized into recurrent "islands" of clinical intent. The mathematical robustness of this partitioning is confirmed by the Silhouette coefficients across the cohort (Fig. 12), which report a consistent median of 0.6. This high degree of separability is visually demonstrated for patient `P01` in Figure 5, where the UMAP projection reveals the convergence of weight-perturbed plans into isolated geometric families.

Quantitative analysis of the $N = 20$ cohort reveals a remarkably stable and balanced distribution of these strategies (Fig. 11). On average, the candidate pool is divided almost symmetrically among the four archetypes: Robust (25.39%), Aggressive (25.36%), Balanced (25.09%), and Conservative (24.16%).

Interestingly, while the majority of the cohort naturally converged into four archetypes, a subset of patients exhibited only three distinct clusters. In these specific cases, the anatomical constraints led to a collapse of two traditionally separate trade-off regimes into a single dominant strategy, highlighting the framework's ability to adapt the decision space to the patient's specific geometry.

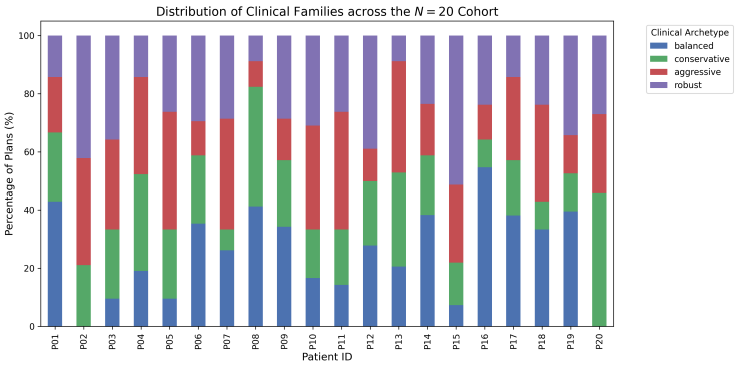

Figure 11: Distribution of clinical archetypes across the $N = 20$ cohort. Each bar representing a single patient shows a consistent and balanced partition of the solution space. The emergence of three clusters in specific cases (P02 and P20) demonstrates the framework's ability to adaptively consolidate redundant trade-off regimes based on individual patient geometry and anatomical constraints.

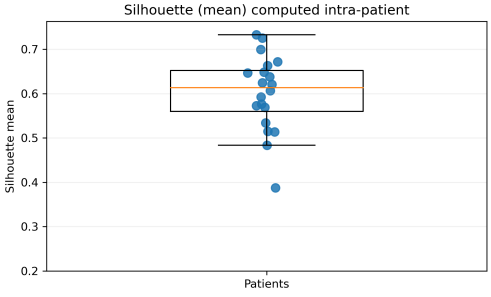

Figure 12: Distribution of intra-patient Silhouette coefficients across the cohort.

The qualitative profiles of these families are summarized in Figures 13 and 14. The signature heatmap exposes sharply differentiated "fingerprints" for each family: *Aggressive* plans prioritize target coverage at the expense of OAR sparing, while *Conservative* plans focus on healthy tissue protection. The transition between these profiles represents a tangible shift in treatment philosophy that far exceeds the intra-cluster variance.

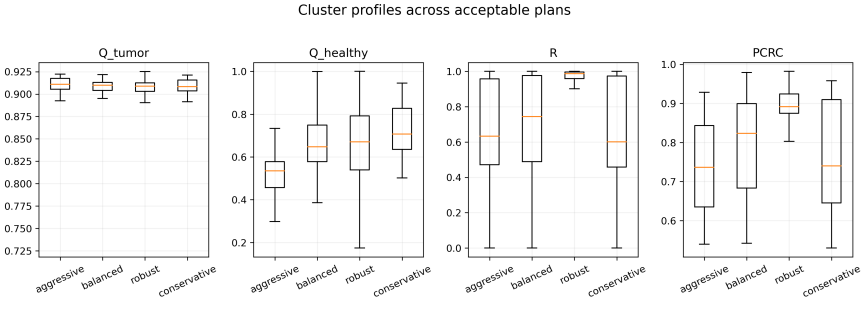

Figure 13: Distributions of quality and robustness metrics across plan families.

Finally, each cluster is summarized by a single representative plan selected using a medoid criterion, i.e., the solution that minimizes the average distance to the other members of the cluster. By delivering a compact subset of 3 to 4 plans, the framework effectively eliminates “choice overload” and provides a principled basis for clinical decision-making. The outcome is a transparent decision support tool that reflects the full spectrum of optimal treatment philosophies for each specific patient.

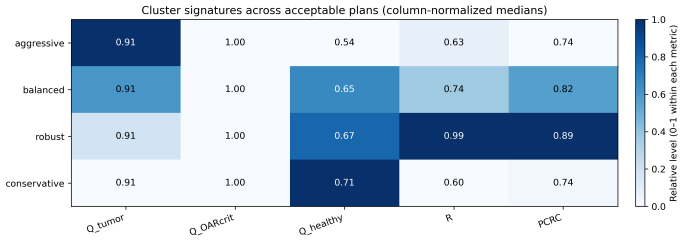

Figure 14: Column-normalized median signatures (fingerprints) of each clinical family.

6. Conclusions

This work introduced a decision-support framework for multi-objective optimization that shifts practice from manual weight tuning toward a structured exploration of the attainable trade-off space under preference uncertainty. By integrating extended Goal Programming with systematic preference perturbations and solution-set structuring, the approach decouples mathematical optimization from subjective preference elicitation and converts large ensembles of admissible solutions into compact, robustness-aware decision sets. The IMPT case study demonstrated the practical viability of this paradigm, showing that complex radiotherapy trade-off landscapes can be compressed into a small number of clinically interpretable treatment strategies.

The main contribution lies in revealing and exploiting the intrinsic redundancy of high-dimensional multi-objective solution spaces. In general, large sets of non-dominated solutions were observed to organize into a limited number of recurrent compromise regimes, indicating that much of the apparent Pareto-set complexity is structurally redundant and can be reduced without loss of decision-relevant information. In the IMPT application, this latent structure manifested as hundreds of treatment plans consistently clustering into a few clinically interpretable strategies across patients. This regularity supports a shift from selecting a single Pareto-optimal solution toward identifying stable regions of the frontier under preference variation. Compared with conventional weighted-sum planning, the framework more reliably exposes organ-sparing opportunities under equivalent target coverage while providing explicit robustness information with respect to goal priorities. By characterizing representative strategies as interpretable archetypes, the methodology also promotes reproducibility and transparency in planning workflows.

These advantages come with important limitations. While the framework explicitly characterizes uncertainty in decision-maker preferences, it does not incorporate uncertainty in underlying problem data or system parameters, which may be critical in many multi-objective optimization settings (including the IMPT application considered here). In addition, the reliance on repeated mixed-integer EGP solves combined with post-optimization clustering entails a higher computational burden than conventional single-solution optimization workflows. The proposed methodology should therefore be interpreted as a decision-support layer that complements, rather than replaces, existing robust optimization approaches, providing structured preference exploration and solution-set interpretation on top of established optimization models.

Within the broader context of multi-criteria decision-making, the proposed approach contributes to the class of *a posteriori* preference modeling methods by integrating systematic sensitivity exploration, robustness characterization, and solution-set reduction in a unified pipeline. In application domains such as proton therapy planning, where most systems rely on *a priori* weight selection or interactive navigation, this perspective shifts the focus from identifying a single preferred solution to characterizing regions of the Pareto frontier that remain stable under preference variation. More generally, the methodology aligns domain-specific planning problems with emerging decision-support paradigms in operations research that emphasize robust regions of trade-offs rather than isolated Pareto-optimal points.

Several directions for future research follow naturally. Extending the framework to incorporate uncertainty in underlying problem data or system parameters would enable joint characterization of preference and operational robustness, including, for example, physical delivery uncertainties in proton therapy planning. Computational acceleration through surrogate models or warm-start strategies could facilitate deployment in time-sensitive decision environments. More broadly, learning mappings between problem characteristics and emergent trade-off regimes may allow rapid prediction of representative solution families. Finally, applying the framework across diverse multi-criteria decision domains will further test its generality as a scalable approach for structuring and interpreting large Pareto solution spaces.

Overall, the results indicate that complex multi-objective optimization landscapes possess exploitable low-dimensional structure that can be systematically organized into interpretable and robustness-aware decision alternatives, providing a reproducible pathway from optimization outputs to actionable decision support across application domains.

Acknowledgments

This study was funded by the Coordination for the Improvement of Higher Education Personnel (CAPES) and the National Council for Scientific and Technological Development – Brazil (CNPq). The authors gratefully acknowledge the support and academic environment of University of Portsmouth during the visiting research period.

CRediT authorship contribution statement

Nicole C. Cassimiro: Conceptualization, Methodology, Software, Formal analysis, Investigation, Data curation, Visualization, Writing - Original draft. **Dylan Jones:** Conceptualization, Methodology, Supervision, Writing - Review & Editing. **Juliana C. Freitas:** Conceptualization, Validation, Investigation, Formal analysis, Supervision, Writing - Review & Editing. **Aurelio R. L. Oliveira:** Supervision, Software, Methodology, Writing - Review & Editing.

References

An, Q., Zhao, J., Zhu, J., 2025. Data envelopment analysis clustering: an optimization based machine learning. Journal of the Operational Research Society , 1–17URL: https://doi.org/10.1080/01605682.2025.2537278, doi:10.1080/01605682.2025.2537278, arXiv:https://doi.org/10.1080/01605682.2025.2537278.

Breedveld, S., Craft, D., van Haveren, R., Heijmen, B., 2017. Multi-criteria optimization and decision-making in radiotherapy. Physics in Medicine & Biology 62, R231–R262. doi:10.1088/1361-6560/aa8f8a.

Breedveld, S., Heijmen, B., 2017. TROTS: The Radiotherapy Optimisation Test Set. URL: https://doi.org/10.5281/zenodo.2708302, doi:10.5281/zenodo.2708302. public benchmark dataset for radiotherapy treatment planning research.

Cao, H., Romeijn, H.E., Shepard, D.M., Craft, D., 2012. An efficient approach to multi-criteria optimization for intensity-modulated radiation therapy treatment planning. Medical Physics 39, 7635–7647. doi:10.1118/1.4762565.

Cao, R., Li, X., Chen, W., Wang, C., Si, L., Pei, X., Zhang, X., 2024. An adaptive conjugate gradient accelerated evolutionary algorithm for multi-objective spot optimization in cancer intensity modulated proton therapy. Applied Soft Computing 151, 111177. URL: https://www.sciencedirect.com/science/article/pii/S156849462301195X, doi:https://doi.org/10.1016/j.asoc.2023.111177.

Corrente, S., Greco, S., Matarazzo, B., Słowiński, R., 2024. Explainable interactive evolutionary multiobjective optimization. Omega 122, 102925. URL: https://www.sciencedirect.com/science/article/pii/S0305048323000890, doi:https://doi.org/10.1016/j.omega.2023.102925.

Corrente, S., Greco, S., Słowiński, R., Zappalà, S., 2026. An explainable and interpretable composite indicator based on decision rules. Omega 142, 103513. URL: https://www.sciencedirect.com/science/article/pii/S0305048326000022, doi:https://doi.org/10.1016/j.omega.2026.103513.

Craft, D., Halabi, T., Shih, H., Bortfeld, T., 2012. Multicriteria optimization for radiation therapy treatment planning. International Journal of Radiation Oncology Biology Physics 82, e39–e45. doi:10.1016/j.ijrobp.2010.11.038.

de Freitas, J.C., Cantane, D.R., Rocha, H., Dias, J., 2024. A multiobjective beam angle optimization framework for intensity-modulated radiation therapy. European Journal of Operational Research 318, 286–296. URL: https://ideas.repec.org/a/eee/ejores/v318y2024i1p286-296.html, doi:10.1016/j.ejor.2024.05.004.

Greco, S., Słowiński, R., Wallenius, J., 2025. Fifty years of multiple criteria decision analysis: From classical methods to robust ordinal regression. European Journal of Operational Research 323, 351–377. URL: https://www.sciencedirect.com/science/article/pii/S0377221724005988, doi:https://doi.org/10.1016/j.ejor.2024.07.038.

Hartikainen, M., Miettinen, K., Klamroth, K., 2019. Interactive nonconvex pareto navigator for multiobjective optimization. European Journal of Operational Research 275, 238–251. URL: https://www.sciencedirect.com/science/article/pii/S0377221718309792, doi:https://doi.org/10.1016/j.ejor.2018.11.038.

Hodapp, N., 2012. Der icru-report 83: Verordnung, dokumentation und kommunikation der fluenzmodulierten photonenstrahlentherapie (imrt). Strahlentherapie und Onkologie 188, 97–100. URL: https://doi.org/10.1007/s00066-011-0015-x, doi:10.1007/s00066-011-0015-x.

Huang, P., Yan, H., Song, Z., Xu, Y., Hu, Z., Dai, J., 2023. Combining autoencoder with clustering analysis for anomaly detection in radiotherapy treatment plans. Quantitative Imaging in Medicine and Surgery 13, 2328–2338. doi:10.21037/qims-22-825.

Huiskes, M., Kong, W., Oud, M., Crama, K., Rasch, C., Breedveld, S., Heijmen, B., Astreinidou, E., 2024. Validation of fully automated robust multicriterial treatment planning for head and neck cancer impt. International Journal of Radiation Oncology*Biology*Physics 119, 968–977. URL: https://www.sciencedirect.com/science/article/pii/S0360301623083086, doi:https://doi.org/10.1016/j.ijrobp.2023.12.034.

International Commission on Radiation Units and Measurements, 2007. Prescribing, recording, and reporting proton-beam therapy. Journal of the ICRU 7, NP–NP. doi:10.1093/jicru/ndm021.

Jones, D., 2011. A practical weight sensitivity algorithm for goal and multiple objective programming. European Journal of Operational Research 213, 238–245. URL: https://www.sciencedirect.com/science/article/pii/S0377221711002232, doi:https://doi.org/10.1016/j.ejor.2011.03.012.

Jones, D., Tamiz, M., 2010. Practical Goal Programming. Springer, New York.

Kamal-Sayed, H., Ma, J., Tseung, H., Abdel-Rehim, A., Herman, M.G., Beltran, C.J., 2018. Adaptive method for multicriteria optimization of intensity-modulated proton therapy. Medical Physics 45, 5643–5652. URL: https://aapm.onlinelibrary.wiley.com/doi/abs/10.1002/mp.13239, doi:https://doi.org/10.1002/mp.13239, arXiv:https://aapm.onlinelibrary.wiley.com/doi/pdf/10.1002/mp.13239.

Li, Z., 2009. Prescribing, recording, and reporting proton-beam therapy. International Journal of Radiation Oncology, Biology, Physics 73, 1602. URL: https://doi.org/10.1016/j.ijrobp.2008.10.084, doi:10.1016/j.ijrobp.2008.10.084.

Lotov, A.V., Bushenkov, V.A., Kamenev, G.K., 2017. Interactive Decision Maps: Approximation and Visualization of Pareto Frontier. Springer, Cham.

Malekpoor, H., Mishra, N., Kumar, S., 2022. A novel topsis–cbr goal programming approach to sustainable healthcare treatment. Annals of Operations Research 312, 1403–1425. URL: https://doi.org/10.1007/s10479-018-2992-y, doi:10.1007/s10479-018-2992-y.

Marler, R.T., Arora, J.S., 2004. Survey of multi-objective optimization methods for engineering. Structural and Multidisciplinary Optimization 26, 369–395. URL: https://doi.org/10.1007/s00158-003-0368-6, doi:10.1007/s00158-003-0368-6.

Merchant, J., 2024. An Investigation into Goal Programming Sensitivity, Duality and Application. Doctoral thesis. University of Portsmouth. URL: https://researchportal.port.ac.uk/en/studentTheses/an-investigation-into-goal-programming-sensitivity-duality-and-ap/.

Mohamed, N., Li, X., Lee, N.Y., 2022. A narrative review of intensity-modulated proton therapy for head and neck cancer. Therapeutic Radiology and Oncology 7. URL: https://tro.amegroups.org/article/view/7610.

Murtagh, F., Contreras, P., 2012. Algorithms for hierarchical clustering: an overview. WIREs Data Mining and Knowledge Discovery 2, 86–97. URL: https://wires.onlinelibrary.wiley.com/doi/abs/10.1002/widm.53, doi:https://doi.org/10.1002/widm.53, arXiv:https://wires.onlinelibrary.wiley.com/doi/pdf/10.1002/widm.53.

Petchrompo, S., Wannakrairot, A., Parlikad, A.K., 2022. Pruning pareto optimal solutions for multi-objective portfolio asset management. European Journal of Operational Research 297, 203–220. URL: https://www.sciencedirect.com/science/article/pii/S037722172100388X, doi:https://doi.org/10.1016/j.ejor.2021.04.053.

Press, R.H., Mehta, M.P., 2024. Proton therapy: Current status and controversies. JCO Oncology Practice 20, 747–749. URL: https://ascopubs.org/doi/abs/10.1200/OP.24.00132, doi:10.1200/OP.24.00132. pMID: 38547434.

Ran, X., Xi, Y., Lu, Y., Wang, X., Lu, Z., 2023. A comprehensive survey of hierarchical clustering algorithms: State-of-the-art machine learning applications, taxonomy, challenges, and future research prospects. Artificial Intelligence Review 56, 8219–8264. doi:10.1007/s10462-022-10249-3.

Romero, C., 2004. Extended lexicographic goal programming: a unifying approach. European Journal of Operational Research 155, 507–517. doi:10.1016/S0377-2217(02)00961-9.

dos Santos Rubem, A.P., Soares de Mello, J.C.C., Angulo Meza, L., 2017. A goal programming approach to solve the multiple criteria dea model. European Journal of Operational Research 260, 134–139. URL: https://www.sciencedirect.com/science/article/pii/S0377221716309742, doi:https://doi.org/10.1016/j.ejor.2016.11.049.

Schubert, M., Teichert, K., 2025. Bi-level multi-criteria optimization to include linear energy transfer into proton treatment planning. Journal of Multi-Criteria Decision Analysis 32. URL: https://onlinelibrary.wiley.com/doi/10.1002/mcda.70021, doi:10.1002/mcda.70021.

Smedberg, H., Bandaru, S., 2023. Interactive knowledge discovery and knowledge visualization for decision support in multi-objective optimization. European Journal of Operational Research 306, 1311–1329. doi:10.1016/j.ejor.2022.09.008.

Sterpin, E., Widesott, L., Poels, K., Hoogeman, M., Korevaar, E.W., Lowe, M., Molinelli, S., Fracchiolla, F., 2024. Robustness evaluation of pencil beam scanning proton therapy treatment planning: A systematic review. Radiotherapy and Oncology 197. URL: https://doi.org/10.1016/j.radonc.2024.110365, doi:10.1016/j.radonc.2024.110365.

Ward, J.H., 1963. Hierarchical grouping to optimize an objective function. Journal of the American Statistical Association 58, 236–244. doi:10.1080/01621459.1963.10500845.

White, S.E., Witing, F., Wittekind, C.I., Volk, M., Strauch, M., 2025. Distilling the pareto optimal front into actionable insights. Environmental Modelling & Software 191, 106508. URL: https://www.sciencedirect.com/science/article/pii/S1364815225001926, doi:https://doi.org/10.1016/j.envsoft.2025.106508.

Wu, Q., Mohan, R., Morris, M., Lauve, A., 2019. Multiobjective optimization in intensity-modulated radiation therapy planning: concepts, methods, and clinical applications. Medical Physics 46, e43–e64. doi:10.1002/mp.13469.

Yadav, D., Nagar, D., Ramu, P., Deb, K., 2023. Visualization-aided multi-criteria decision-making using interpretable self-organizing maps. European Journal of Operational Research 309, 1183–1200. URL: https://www.sciencedirect.com/science/article/pii/S0377221723001145, doi:https://doi.org/10.1016/j.ejor.2023.01.062.

Zhang, P., Fan, N., Shan, J., Schild, S.E., Bues, M., Liu, W., 2017. Mixed integer programming with dose-volume constraints in intensity-modulated proton therapy. Journal of Applied Clinical Medical Physics 18, 29–35. URL: https://aapm.onlinelibrary.wiley.com/doi/abs/10.1002/acm2.12130, doi:https://doi.org/10.1002/acm2.12130, arXiv:https://aapm.onlinelibrary.wiley.com/doi/pdf/10.1002/acm2.12130.