A consolidation heuristic for a biobjective multimodal transportation planning problem

Dominik Leib
Fraunhofer ITWM
Kaiserslautern, RLP, Germany
dominik.leib@itwm.fraunhofer.de

Neele Leithäuser
Fraunhofer ITWM
Kaiserslautern, RLP, Germany
neele.leithaeuser@itwm.fraunhofer.de

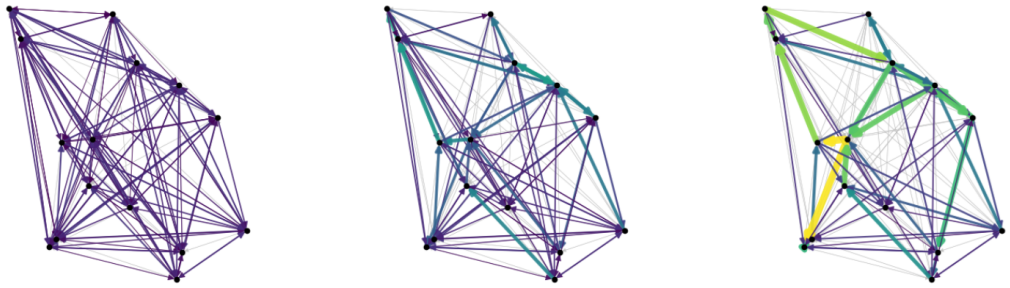

Figure 1: Three passenger flows of the same volume of varying degree of consolidation.

Abstract

This paper addresses multimodal transportation planning with an explicit focus on multiple modes and two objectives. While traditional approaches typically optimize single modes and emphasize user-oriented criteria such as travel time and convenience, environmental impacts—especially carbon emissions and energy consumption—are becoming equally important. We therefore propose a bi-criteria optimization model for multimodal network design that captures the trade-off between user attractiveness and energy consumption. To manage the computational complexity of the resulting problem, we develop a heuristic that restricts the lower-level search space to a small set of promising flow candidates. Using this framework, we compute Pareto frontiers for illustrative case studies and analyze the resulting trade-offs. The findings show how different priorities between travel efficiency and energy reduction lead to distinct multimodal network configurations, providing a basis for more transparent, trade-off-aware planning decisions.

Keywords

bi-objective optimization, heuristic methods, flow consolidation, Pareto frontier approximation

1 Motivation

Multimodal transportation planning focuses on designing and managing travel systems that integrate multiple modes—such as walking, cycling, public transit, and private vehicles—into a coherent whole. Rather than optimizing each mode in isolation, multimodal planning evaluates how different modes interact and complement each other across entire trips. This integrated perspective is increasingly important in urban and regional contexts, where capacity constraints, congestion, and changing mobility needs demand flexible, interconnected solutions rather than single-mode expansions. For a recent review on modelling approaches, see [?].

A key challenge in multimodal planning is that it is inherently multi-criteria. Traditional transport planning often prioritizes travel time and user convenience, aiming to minimize delays and improve reliability. However, growing concerns about climate change require that environmental criteria, especially carbon emissions and energy consumption, are treated as central objectives of equal importance. This creates trade-offs: the fastest or most convenient option for individuals may not be the lowest-energy option for society. Multimodal, multi-criteria planning therefore seeks to balance travel time, comfort, accessibility, and cost with energy consumption and other sustainability indicators, enabling decision-makers to evaluate alternatives in a more holistic and transparent manner. The bi-criteria planning problem of minimizing emission and travel-time has been investigated in e. g. [?] for the unimodal case.

We present a bi-criteria optimization model for multimodal network design that jointly considers user-oriented performance (e.g., travel time and convenience) and energy consumption, while not explicitly accounting for user costs and infrastructure investment costs. To handle the resulting complexity, we propose a heuristic that restricts the continuous flow search space to a small set of promising candidates. Based on this approach, we compute and analyze Pareto frontiers for illustrative multimodal use cases.

2 The Modal Split Model

In this section we introduce a baseline bi-objective modal split network design model. We consider a directed graph G together with a finite set of transport modes M. Each edge may be operated in each mode, which induces mode-dependent travel time and energy consumption. Passenger demand is given by a set of origin–destination pairs (commodities) with a travel demand. The demand may be distributed across available modes and routes (‘modal split’), yielding a multi-commodity flow on the network.

INOC ’26, Liège (Belgium)

10.48786/inoc.2026.21

The model jointly decides mode-specific service levels on edges and how passenger flow is routed on the resulting infrastructure and falls into the class of bilevel network design problems.

Sets and indices.

- V: set of nodes.
- E: set of directed links, $e = (u, v) \in E$.
- M^{PT}: set of public transport modes (designable modes).
- m^{ind}: individual transport, always available.
- $M := M^{\mathrm{PT}} \cup \{m^{\mathrm{ind}}\}$: set of all modes.
- $\mathcal{D} \subseteq V$: set of destinations.

Parameters.

- $D_{o,d} \geq 0$: demand from origin o to destination d, the pairs (o, d) with $D_{o,d} > 0$ are the commodities K
- $\ell_e \geq 0$: length (distance) of link e.
- $\tau_m > 0$: average time of mode $m \in M$ (time per unit of distance)
- $\eta_m > 0$: average energy consumption per vehicle and unit of distance for $m \in M$.
- $\mathrm{CAP}_m > 0$: passenger capacity provided by one vehicle of public transport mode $m \in M^{\mathrm{PT}}$.

Decision variables.

- $x_{d,e,m} \geq 0$: passenger flow with destination d on link e in mode m (destination-based flow), for all $m \in M$.
- $y_{e,m} \in \mathbb{Z}_+$: number of vehicles of public transport mode $m \in M^{\mathrm{PT}}$ operated on link e.

Individual vehicle count (derived). For the individual transport we assume unit capacity and identify vehicle-kilometers with passenger-kilometers, implying continuity of this mode. Accordingly, we define the induced number of individual transport vehicles on edge e as

$$y_{e,\mathrm{ind}} := \sum_{d \in \mathcal{D}} x_{d,e,\mathrm{ind}},$$

which implicitly sets $\mathrm{CAP}_{\mathrm{ind}} = 1$.

Flow conservation. The model is destination-based: for each destination $d \in \mathcal{D}$ we route all demands with destination d through the network. Demand may split across modes and routes and we do not penalize intermodal transfers in this baseline.

For all $d \in \mathcal{D}$ and all $w \in V \setminus \{d\}$:

$$\sum_{m \in M} \sum_{e=(w,u) \in E} x_{d,e,m} - \sum_{m \in M} \sum_{e=(u,w) \in E} x_{d,e,m} = D_{w,d}. \tag{1}$$

At the destination d we require that no flow leaves:

$$\sum_{m \in M} \sum_{e=(d,u) \in E} x_{d,e,m} = 0. \tag{2}$$

Then (1) implies that the total inflow into d equals $\sum_{w \in V} D_{w,d}$.

Capacity and design constraints. Passenger flows are linked to operated capacity. For all $e \in E$, $m \in M^{\mathrm{PT}}$:

$$\sum_{d \in \mathcal{D}} x_{d,e,m} \leq \mathrm{CAP}_m \, y_{e,m}, \tag{3}$$

Objectives. We consider two criteria: total in-vehicle travel time and operational energy consumption. Travel time is proportional to link length with mode-dependent speed, and energy consumption is induced by the operated fleet.

$$\mathrm{TIME}(x, y) = \mathrm{TIME}(x) = \sum_{d \in \mathcal{D}} \sum_{e \in E} \sum_{m \in M} \ell_e \tau_m \, x_{d,e,m}, \tag{4}$$

$$\mathrm{ENERGY}(x, y) = \sum_{e \in E} \sum_{m \in M} \eta_m \, \ell_e \, y_{e,m} \tag{5}$$

$$= \sum_{e \in E} \ell_e \left(\sum_{m \in M^{\mathrm{PT}}} \eta_m \, y_{e,m} + \eta_{\mathrm{ind}} \sum_{d \in \mathcal{D}} x_{d,e,\mathrm{ind}} \right) \tag{6}$$

We do not impose explicit upper bounds or a budget on the service levels $y_{e,m}$. Installing additional vehicles increases the energy objective via (6) and is therefore penalized in the bi-objective trade-off. The problem then reads

$$\begin{aligned} \min_{x,\, y} \quad & \big(\mathrm{TIME}(x),\ \mathrm{ENERGY}(x, y)\big) \\ \text{s.t.} \quad & (1), (2),\ (3), \\ & x \geq 0,\ y \geq 0. \end{aligned} \tag{MSM}$$

The corresponding decision problem to MSM ist NP-complete and also hard to approximate, even for just one public mode of transportation in M^{PT}. But with a fixed passenger flow, it is easier to solve and even provides a FPTAS in the latter case (see [**?**] for details). This gives rise to the idea of fixing the passenger paths to some carefully selected flows, which is the baseline for the heuristic introduced in the upcoming sections.

3 Computing Consolidated Flows

We want to precompute mode-independent passenger flows to reduce the lower level search space of the bilevel problem introduced above. The idea is that with energy reduction the passengers consolidate into modes of high capacity and therefore reduce the per head energy consumption. To obtain structured flow patterns quickly, we compute such *consolidated* (passenger) flows by means of a lightweight surrogate objective. Intuitively, the surrogate rewards solutions in which multiple OD-demands share common corridors, which is particularly relevant in dense networks with many near-shortest alternatives.

Passenger flows and aggregated edge usage. Let K be the set of OD-pairs (commodities). For each commodity $k = (s_k, t_k) \in K$ with demand $D_k \geq 0$ we use non-negative flow variables $f_e^k \geq 0$ indicating the amount of passenger demand of k routed on directed edge $e \in E$. The *aggregated* passenger flow on an edge is

$$F_e := \sum_{k \in K} f_e^k.$$

Concave consolidation surrogate. To encourage consolidation, we penalize aggregated edge usage by a concave, increasing function $\phi_\gamma : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0}$, e.g., $\phi_\gamma(u) = u^\gamma$ with $\gamma \in (0, 1)$. Using the edge lengths ℓ_e, the surrogate objective takes the form

$$J_\gamma(F) := \sum_{e \in E} \ell_e \, \phi_\gamma(F_e). \tag{7}$$

Concavity of ϕ_γ induces economies of scale and thus favors routing multiple commodities through shared corridors.

Note that this approach is inspired by [**?**] and [**?**], but has the opposite objective and therefore a disjoint choice of γ. Instead of distributing individual traffic in order to eliminate road congestion, we seek to bundle the transport demand in order to route them with public transport vehicles.

Iterative reweighting. Directly minimizing (7) is non-convex, in fact, the flow problem listed above is NP-hard. We therefore apply an iterative reweighting scheme that repeatedly linearizes the concave term. Given the current aggregated usage $F^{(r)}$ in iteration r, we define edge weights

$$w_e^{(r)} := \ell_e \cdot \gamma \cdot \left(F_e^{(r)} + \varepsilon\right)^{\gamma-1}, \tag{8}$$

where $\varepsilon > 0$ is a small stabilization constant preventing infinite weights at $F_e^{(r)} = 0$.

In iteration $r+1$, we route each commodity k on a shortest path from s_k to t_k with respect to the weights $w^{(r)}$ and assign its full demand D_k to that path (all-or-nothing assignment). This produces updated flows $f^{(r+1)}$ and aggregated usage $F^{(r+1)}$. We iterate until the relative improvement of (7) falls below a tolerance or a maximum number of iterations is reached; see Algorithm 1. Afterwards we aggregate by destination to achieve destination based flows $f_\gamma := (f^d)_{d\in\mathcal{D}} = (\sum_o f^{(o,d)})_{d\in\mathcal{D}}$.

Algorithm 1 Iteratively reweighted shortest paths for consolidated OD flows

Require: Directed graph $G = (V, E)$ with lengths ℓ_e; OD demands $(s_k, t_k, D_k)_{k\in K}$; parameter $\gamma \in (0, 1)$; $\varepsilon > 0$; max iterations R; tolerance Tol_{IRSP}.
Ensure: OD flow pattern $(f^k)_{k\in K}$

Initialize $F_e^{(0)} \leftarrow 0$ for all $e \in E$
$J^{(0)} \leftarrow \sum_{e\in E} \ell_e \left(F_e^{(0)}\right)^\gamma$
for $r = 0, 1, 2, \ldots, R-1$ **do**
  Set $w_e^{(r)} \leftarrow \ell_e \cdot \gamma \cdot \left(F_e^{(r)} + \varepsilon\right)^{\gamma-1}$ for all $e \in E$
  Initialize $f_e^k \leftarrow 0$ for all $k \in K$, $e \in E$
  for all $k \in K$ **do**
    Compute a shortest path P_k from s_k to t_k w.r.t. $w^{(r)}$
    for all $e \in P_k$ **do**
      $f_e^k \leftarrow D_k$
    end for
  end for
  $F_e^{(r+1)} \leftarrow \sum_{k\in K} f_e^k$ for all $e \in E$
  $J^{(r+1)} \leftarrow \sum_{e\in E} \ell_e \left(F_e^{(r+1)}\right)^\gamma$
  if $\dfrac{J^{(r)} - J^{(r+1)}}{\max\{1, J^{(r)}\}} < \text{Tol}_{IRSP}$ **then**
    break
  end if
end for
return $(f^k)_{k\in K}$

Candidate generation. We run the above procedure for a small grid of parameters $\Gamma \subset (0, 1]$ and collect the resulting consolidated OD flows, where the flow for $\gamma = 1$ is simply obtained by a all-pairs shortest path algorithm. The union over $\gamma \in \Gamma$ provides a compact set of candidate flow patterns that interpolates between near-shortest routing (weak consolidation and large γ) and strongly shared corridors (strong consolidation and small γ). These candidates are then evaluated in terms of the original bi-objective criteria introduced in Section 2. One can show that in general there are, up to symmetry, only finitely many optimal solutions to the flow problem, which are furthermore optimal for disjoint subintervals of $(0, 1]$. This may imply duplicates while sampling, which will be omitted. Figure 1 exemplarily shows the flows of the complete random graph with high demand use case (cf. Section 5) for $\gamma \in \{1.0, 0.85, 0.71\}$.

4 Methodology

Patch Approximation via Sandwiching. Our baseline model (MSM) is a bi-objective mixed-integer program due to the public-transport service decisions $y_{e,m}$ ($m \in M^{\text{PT}}$). Fixing all discrete design variables to a given *layout* (i.e., fixing the PT service levels $y_{e,m}$) yields a continuous bi-objective problem in the remaining flow variables x. In our setting, the resulting subproblem is a bi-objective linear program; hence its Pareto frontier in the objective space is convex and consists of piecewise-linear segments. We call those subproblems *patch problems*, whereas their Pareto frontier is called a *patch*. As a consequence, the Pareto frontier of the integrated problem MSM is then the non dominated set of all patches.

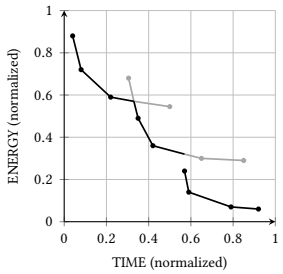

Figure 2: Pareto frontier of integrated problem in black of three piecewise-linear convex patches.

We approximate this frontier by a Sandwiching scheme (see e. g. [**? ?**]): starting from the two single-objective extrema (weights $(1, 0)$ and $(0, 1)$), we iteratively solve weighted-sum scalarizations with weights $\lambda \in \mathbb{R}^2_{>0}$. Each solve yields a supported non dominated point and a corresponding supporting hyperplane in objective space. The convex hull of the discovered points (augmented by the domination cone) forms an inner approximation, while the intersection of the supporting half-spaces forms an outer approximation.

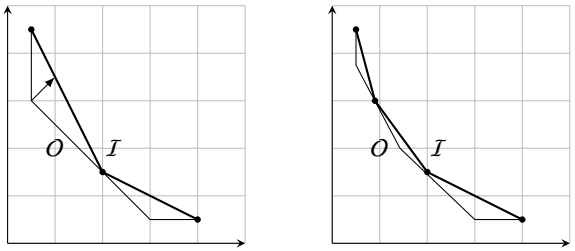

Figure 3: A sandwiching step: the current approximation is improved at the segment, where the epsilon-indicator measure is maximized.

A key computational ingredient of Sandwiching algorithms is the evaluation of the current approximation quality (i.e., the gap between inner and outer approximations). We employ the standard quality (multiplicative) epsilon-indicator, for which the distance evaluation can be formulated as small linear programs. The epsilon-indicator is used as a metric to identify the point of maximal discrepancy towards $(1, 1)$ direction between inner and outer approximation and choose the next weight accordingly, thereby tightening the outer bound and extending the inner hull until the desired accuracy is reached. Definitions, efficient update

Algorithm 2 Sandwiching for patch approximation

Require: Fixed PT layout y^{PT} and tolerance $\mathrm{Tol}_{SW} > 0$.
Ensure: Supported non dominated point set P approximating the convex patch frontier.

1: Solve the two single-objective extrema (weights $(1,0)$ and $(0,1)$) and set $\mathrm{P} \leftarrow \{p^{(1)}, p^{(2)}\}$
2: Initialize inner approximation $\mathcal{I} \leftarrow \mathrm{conv}(\mathrm{P}) + \mathbb{R}^2_{\geq 0}$
3: Initialize outer approximation O as the intersection of supporting half-spaces induced by P
4: **while** true **do**
5: Compute the current discrepancy Δ (epsilon-indicator) between $\mathcal{I}$ and O
6: **if** $\Delta \leq \mathrm{Tol}_{SW}$ **then**
7: **break**
8: **end if**
9: Select a weight vector $\lambda \in \mathbb{R}^2_{>0}$ corresponding to a maximally violated facet / gap of O
10: Solve weighted-sum scalarization of the patch problem:

$$x^\star \in \arg\min \lambda_1 \,\mathrm{TIME}(x) + \lambda_2 \,\mathrm{ENERGY}(x,y)$$
$$\text{s.t. } x \text{ feasible for fixed } y$$

11: $p \leftarrow \big(\mathrm{TIME}(x^\star), \mathrm{ENERGY}(x^\star, y)\big)$; $\mathrm{P} \leftarrow \mathrm{P} \cup \{p\}$
12: Update $\mathcal{I} \leftarrow \mathrm{conv}(P) + \mathbb{R}^2_{\geq 0}$ and refine O by adding the supporting half-space at p with normal λ
13: **end while**
14: **return** P

criteria and practical speed-ups for these quality computations are described in [**?**].

Equidistant Frontier Sampling by Enumeration Solver. As a baseline, we sample the Pareto frontier of the integrated mixed-integer model (MSM) by an ε-constraint strategy on the TIME objective. We first compute the two Pareto extrema by solving (MSM) once with objective min TIME and once with objective min ENERGY (subject to (1)–(3)). This yields a TIME-range $[\mathrm{TIME}_{\min}, \mathrm{TIME}_{\max}]$.

We then choose a user-defined number of samples $S \geq 2$ and generate equidistant thresholds $\tau_i := \mathrm{TIME}_{\min} + \frac{i}{S-1}\big(\mathrm{TIME}_{\max} - \mathrm{TIME}_{\min}\big)$ for $i = 1, \ldots, S-2$. For each interior threshold $i = 1, \ldots, S-2$, we solve the constrained problem

$$\min \ \mathrm{ENERGY}(x,y) \quad \text{s.t. } (1),(2),(3),\ \mathrm{TIME}(x) \leq \tau_i, \qquad (9)$$

using a prescribed time limit. Each resulting integrated solution induces a fixed layout y^{PT}; for every such layout we subsequently compute the corresponding patch by the Sandwiching-based patch approximation Algorithm 2 up to approximation tolerance Tol_{SW}.

Finally, we take the union of all patches obtained from the sampled layouts (including the extrema), and extract the non dominated subset as our final approximation of the integrated Pareto frontier. The procedure is controlled by three parameters: the number of samples S, the time limit t for the integrated solves, and the patch approximation tolerance Tol_{SW}.

Equidistant Frontier Sampling by Consolidation Heuristic. Our heuristic approach replaces the full integrated search over layouts by a compact candidate set derived from consolidated passenger flows. For a given parameter grid $\Gamma \subset (0,1]$ we compute consolidated OD-flow patterns $\{f_\gamma, \gamma \in \Gamma\}$, using the concave surrogate and IRSP reweighting Algorithm 1. Each f_γ induces a structured corridor usage pattern a feasible (mode-aggregate) flow for MSM.

The flow space in the integrated problem is then reduced to the choices of the computed consolidated flows $\{f_\gamma \mid \gamma \in \Gamma\}$ by additional constraints

$$\sum_{m \in M} x_{d,e,m} = b_\gamma f^d_\gamma(e) \quad \forall e \in E, d \in \mathcal{D}, \quad \sum_{\gamma \in \Gamma} b_\gamma = 1, \qquad (10)$$

where b_γ is a binary variable for each $\gamma \in \Gamma$, ensuring that exactly one consolidated flow is chosen. We then proceed as above and compute the two Pareto extrema of the resulting restricted mixed-integer model and generate S equidistant TIME-thresholds. For each threshold we solve an ε-constraint subproblem of the form

$$\min \ \mathrm{ENERGY}(x,y) \quad \text{s.t. } (1),(2),(3),(10),\ \mathrm{TIME}(x) \leq \tau_i \quad (11)$$

up to a prescribed time limit. As before, each obtained layout is subsequently refined by computing its continuous patch up to the tolerance. Finally, we take the union over all candidates $\gamma \in \Gamma$ and all sampled thresholds and extract the non dominated subset as the heuristic approximation of the Pareto frontier.

In summary, compared to exact equidistant enumeration, the heuristic differs only by restricting the integrated model to a smaller, consolidation-induced subspace; this typically reduces the number and difficulty of MIP solving while preserving the overall sampling and patch-refinement pipeline.

5 Experimental Results

Setup. For the benchmark we do experiments on two directed, symmetric graphs, the Mandl graph [**?**] and a random complete graph with 15 nodes each, both are shown below. The random graph nodes are sampled in $[0,10]^2$ with euclidean lengths, whereas the Mandl graph is a well known benchmark graph in the transport planning literature and is non-euclidean.

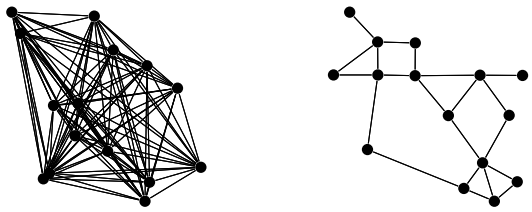

Figure 4: The two input graphs with $n = 15$ each, a directed, symmetric, random complete graph on the left, and the (directed, symmetric) Mandl network on the right.

For each graph we generate a random demand D with parameters v, p, where v is the maximal demand per commodity, and p is the probability that the demand of one commodity is non-zero. We consider per graph two use cases, a low demand with $(15, 60\%)$ and a high demand with parameters $(75, 80\%)$ to investigate effects of varying demand on the consolidation strategy. Two discrete modes are available in PT, a small one (bus) and a large one (tram). Parameters were derived from practical data and are listed in the following table, besides the same parameter set for individual transport (ind).

We give the solver an upper time limit of 600s and compute 6 samples per use case. We chose Gurobi as a solver on a machine with an Intel i7-1370P with 14 Cores and 32GB memory. The Sandwiching algorithm is then applied to each sample to achieve one patch each, where we set the tolerance for the patch approximation to $\mathrm{Tol}_{SW} = 0.01$. In each plot we only show the non

Mode	Capacity [Pers.]	Energy Consumption [kWh/km/vehicle]	Time [h/km]
ind	1	0,20	0,025000
bus	42	1,15	0,036236
tram	164	4,70	0,028519

Table 1: Mode Vehicle Parameters

dominated parts, therefore some patches might vanish. For the flow generation we set $\Gamma := \{i\frac{1}{15} \mid i = 1, \ldots, 15\}$ as a sampling set for the consolidation factors. We set the max iterations to $R = 100$, the tolerance to $\text{Tol}_{IRSP} = 10^{-7}$ and $\varepsilon := 10^{-6}$. In each plot we show both the frontier of the baseline in black and the frontier of the heuristic in red in a (unified) normalized scale. Additionally the plots contain the γ values corresponding to the flows chosen by the heuristic, which are close to the patch they correspond to.

Metrics. We compare the approximated Pareto frontiers **A** of the baseline with the exact solver and the frontier **B** of the heuristic and use the following notation and indicators, setting $\mathbf{U} = \text{ND}(\mathbf{A} \cup \mathbf{B})$ to be the the non-dominated subset of the union of both frontiers as the best known approximation of the true Pareto frontier.

- **Hypervolume difference.** We normalize objectives to the bounding box of **U** and set the reference point r to its upper-right corner. For $\mathbf{F} \in \{\mathbf{A}, \mathbf{B}, \mathbf{U}\}$, $\text{HV}(\mathbf{F})$ denotes the dominated area with respect to r in the normalized space. We report
$$\Delta\text{HV} = \text{HV}(\mathbf{A}) - \text{HV}(\mathbf{B}),$$
where a larger value indicates that **A** dominates a larger portion of the objective space than **B**.
- **Additive epsilon-indicator.** The epsilon-indicator $I_\varepsilon(\mathbf{A}, \mathbf{B})$ is defined as
$$I_\varepsilon(\mathbf{A}, \mathbf{B}) = \inf \{\varepsilon \in \mathbb{R} \mid \forall b \in \mathbf{B}\ \exists a \in \mathbf{A} : a_i \le b_i + \varepsilon \ \ \forall i\}.$$
Smaller values mean that **A** is at least as good as **B** (or better) in all objectives, up to a small additive shift ε towards $(1, 1)$.
- **Distance to the union frontier U.** Let $\widehat{\mathbf{U}}, \widehat{\mathbf{F}}$ be a set of sample points obtained by sampling the polylines of **U** and **F**. We report the (inverted) generational distance
$$I_{\text{GD}}(\mathbf{F}) = \frac{1}{|\widehat{\mathbf{U}}|} \sum_{u \in \widehat{\mathbf{U}}} \min_{p \in \widehat{\mathbf{F}}} \|p - u\|_2.$$
Smaller values indicate that **F** lies closer (on average) to the best-known frontier **U** under the chosen sampling resolution, chosen as 500 points per sub segment.

All metrics are computed in the normalized objective space, values below 0.005 are reported as ~ 0. Additionally we show the runtimes $t(\mathbf{A}), t(\mathbf{B})$ in seconds for computing both frontiers for each use case. For the heuristic, this also includes the consolidation flow generation. The results for the random graph with large and small demand are illustrated in Figure 5 and Figure 6, the ones for the Mandl graph in Figure 7 and Figure 8, respectively.

Overall, the solution quality is consistently strong. While the low-demand use cases exhibit a slightly larger deviation—the heuristic still produces Pareto frontiers that remain close to those obtained by the baseline in all instances.

Across all use cases, the consolidation heuristic is also markedly faster than the baseline, producing speedups of about 2.9 times up

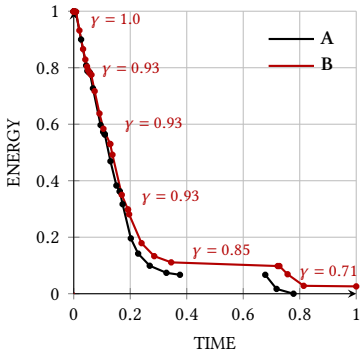

Metric	Val.
ΔHV	0.04
$I_\varepsilon(\mathbf{A}, \mathbf{B})$	~ 0
$I_\varepsilon(\mathbf{B}, \mathbf{A})$	0.05
$I_{\text{GD}}(\mathbf{A})$	~ 0
$I_{\text{GD}}(\mathbf{B})$	0.02
$t(\mathbf{A})[s]$	624.4
$t(\mathbf{B})[s]$	42.5

Figure 5: Use case (random, high demand): normalized (TIME, ENERGY) patch frontiers, exact enumeration vs. heuristic.

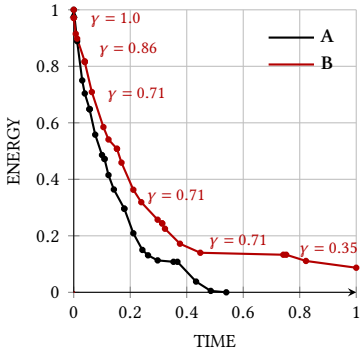

Metric	Val.
ΔHV	0.12
$I_\varepsilon(\mathbf{A}, \mathbf{B})$	~ 0
$I_\varepsilon(\mathbf{B}, \mathbf{A})$	0.13
$I_{\text{GD}}(\mathbf{A})$	~ 0
$I_{\text{GD}}(\mathbf{B})$	0.06
$t(\mathbf{A})[s]$	1112.7
$t(\mathbf{B})[s]$	17.8

Figure 6: Use case (random, low demand): normalized (TIME, ENERGY) patch frontiers, exact enumeration vs. heuristic.

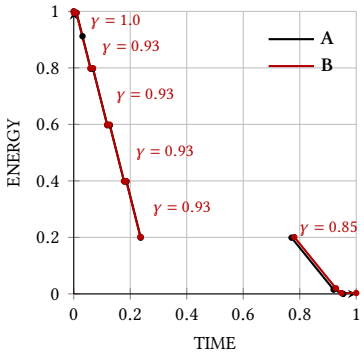

Metric	Val.
ΔHV	~ 0
$I_\varepsilon(\mathbf{A}, \mathbf{B})$	~ 0
$I_\varepsilon(\mathbf{B}, \mathbf{A})$	0.01
$I_{\text{GD}}(\mathbf{A})$	~ 0
$I_{\text{GD}}(\mathbf{B})$	~ 0
$t(\mathbf{A})[s]$	905.9
$t(\mathbf{B})[s]$	315.79

Figure 7: Use case (Mandl, high demand): normalized (TIME, ENERGY) patch frontiers, exact enumeration vs. heuristic.

to over 60 times. The particularly low runtimes on low demand use cases also reflect their discrete nature, as with fewer competitive routing and mode alternatives, the search space collapses quickly and the heuristic converges after only a small number of meaningful design changes.

Lastly we consistently observe that the γ values associated with heuristic patches decrease as ENERGY decreases along the Pareto frontier. Candidates that attain lower energy levels typically correspond to more strongly consolidated corridor flows,

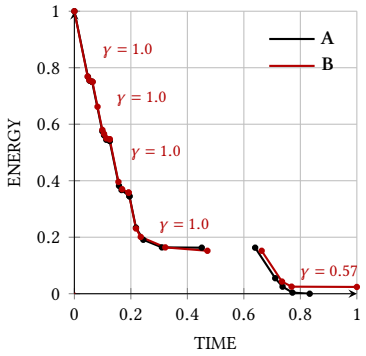

Metric	Val.
ΔHV	0.01
$I_\varepsilon(\mathbf{A}, \mathbf{B})$	0.01
$I_\varepsilon(\mathbf{B}, \mathbf{A})$	0.02
$I_{\text{GD}}(\mathbf{A})$	~ 0
$I_{\text{GD}}(\mathbf{B})$	~ 0
$t(\mathbf{A})[s]$	100.3
$t(\mathbf{B})[s]$	3.3

Figure 8: Use case (Mandl, low demand): normalized (TIME, ENERGY) patch frontiers, exact enumeration vs. heuristic.

which in turn induce higher utilization of public transport infrastructure and reduce individual transport in the resulting modal split. This aligns with the intended role of the surrogate objective as a fast generator of flows that provoke solutions of low energy consumption in the solving process.

6 Conclusion and Outlook

We presented a consolidation-based candidate generation heuristic for a bi-objective modal split network design setting. The method computes corridor-like OD flow patterns via a concave surrogate and iteratively reweighted shortest paths, and evaluates resulting layouts through the same patch-approximation routine as the exact enumeration. Across our test cases, the heuristic reproduces the qualitative front structure while substantially reducing runtime, mainly by reduction of the search space. Future work includes incorporating mode-dependent information into the bundling surrogate, refining the IRSP solver, and evaluating scalability on larger networks and broader demand regimes.

Acknowledgments

This work has been funded by the Federal Ministry of Research, Technology and Space (BMFTR) under the project "Synergien aus physikalischen und verkehrsplanerischen Modellen zur multikriteriellen Optimierung multimodaler nachfrageorientierter Verkehre" (05M22AMC). The authors are grateful for the support. We like to thank the project partners Karl-Heinz Küfer, Anita Schöbel, Sven Jäger, Sarah Roth and Lena Dittrich for useful discussions regarding the model, Markus Friedrich as well as Julian Zimmer for the vehicle parameters and Michael Markl and Tobias Harks for further discussions.