

# AlphaTransit: Learning to Design City-scale Transit Routes

Bibek Poudel<sup>1</sup> Sai Swaminathan<sup>1</sup> Weizi Li<sup>2</sup>

<sup>1</sup>Department of EECS, University of Tennessee, Knoxville, TN, USA

<sup>2</sup>Department of CSE, University of California, Riverside, CA, USA

Correspondence: bpoudel13@vols.utk.edu




Figure 1: Learning dynamics and search cost under mixed demand ( $\alpha = 0.3$ ). **LEFT:** Curves show averages over two training seeds per reward mode. Reward shaping is critical for End-to-End RL. Early Stopping (ES) penalizes routes that terminate before  $L_{\max}$ , while Delta-coverage ( $\Delta$ ) rewards only newly covered demand. The  $\Delta$  variants are strongest;  $\Delta + \text{No ES}$  is best. **MIDDLE:** MCTS search improves sample efficiency under the same environment-step budget. At search depth  $N_{\text{iter}} = 500$ , AlphaTransit surpasses End-to-End RL around  $3 \times 10^5$  steps and finishes with a final smoothed reward 1.31 versus  $-4.02$ . **RIGHT:** Learned policy and value estimates make MCTS search practical across search depths  $N_{\text{iter}} \in [100, 500]$ , when benchmarked on the Bloomington network with a single CPU worker. AlphaTransit stays within seconds per decision, whereas Pure MCTS requires hundreds to thousands of seconds per decision.

## Abstract

Designing a transit network requires many sequential route extension decisions, but their quality is often visible only after the full network is assembled. This delayed-feedback challenge lies at the heart of the Transit Route Network Design Problem (TRNDP), where route interactions can be deceptive: an extension that appears useful locally can create transfer bottlenecks, produce redundant overlap, or reduce overall throughput. To guide route construction under delayed simulator feedback, we introduce AlphaTransit, a search-based planning framework for city-scale bus network design. AlphaTransit couples Monte Carlo Tree Search (MCTS) with a neural policy-value network: the policy proposes route extensions, the value estimates downstream design quality, and search uses these predictions to refine each decision. This provides decision-time lookahead during route construction without running simulator rollouts inside the search tree. We evaluate AlphaTransit on a new Bloomington TRNDP benchmark with realistic road topology and census-derived demand, under mixed and full transit demand settings. In the Bloomington network, AlphaTransit attains the highest service rate in both demand settings, reaching 54.6% and 82.1%, respectively. Relative to reinforcement learning without search, these correspond to 9.9% and 11.4% service rate gains; relative to MCTS without learned guidance, they correspond to 2.5% and 11.2% gains. These results suggest that coupling learned guidance with MCTS is more effective than




Figure 2: AlphaTransit overview. At each route-construction state  $s_t$ , MCTS uses the policy-value network  $f_\theta = (P_\theta, V_\theta)$  to perform selection, expansion, evaluation, and backpropagation. The visit count statistics define a MCTS policy  $\pi_t$ , from which the next action is sampled to advance the state. After the full route set is completed, UXsim evaluates the design and returns a terminal reward  $z$ . Tuples  $(s_t, \pi_t, z)$  are stored in buffer  $\mathcal{D}$  and sampled to train  $f_\theta$  with policy and value losses.

using either approach alone for transit network design. Our code and data are publicly available in <https://github.com/poudel-bibek/AlphaTransit>.

## 1 Introduction

Learning to act under sparse rewards is a central challenge in sequential decision making [54]. In such settings, combining policy priors and value estimates with Monte Carlo Tree Search (MCTS) has notably achieved superhuman performance in games such as Go, chess, and shogi [51, 53, 52]. This approach has also succeeded with unknown dynamics by planning with a learned model [47], and has yielded breakthroughs in algorithmic discovery for matrix multiplication [19] and low-level sorting routines [35]. Together, these results suggest that learned priors combined with explicit lookahead extract a stronger training signal from sparse reward environments than policy learning alone.

We study this principle in *Transit Route Network Design Problem (TRNDP)*, a combinatorial optimization problem in which a planner designs a set of possibly overlapping routes on a shared street network to serve many-to-many passenger flows under operational constraints such as fleet size and cost [46, 17]. Unlike classical routing problems such as the travelling salesman problem [31] or the vehicle routing problem [55], TRNDP jointly designs a route network whose components interact through transfers, congestion, and shared infrastructure [11]. As a result, changing one route can reassign passengers across the network and alter the value of distant route segments. The reward signal is therefore delayed and nonlocal, since the effect of each route extension is observed only after the full network is assembled and simulated. TRNDP is NP-hard [28, 43]; the Bloomington setting studied in this work admits  $\approx 10^{82}$  candidate route sets (Appendix A). The optimization landscape is also deceptive: a locally attractive route extension can create transfer bottlenecks, redundant overlap, and lower global throughput. In addition, the objective is inherently multi criteria, balancing passenger facing metrics such as coverage, waiting time, and journey time against operator side constraints such as fleet size, route overlap, and vehicle utilization.

Prior TRNDP methods handle this difficulty by simplifying the evaluation model, iteratively modifying candidate route sets, or learning route construction policies. Exact and decomposition-based methods follow the first path, making optimization tractable by using analytical passenger routing and travel time objectives rather than simulation [5, 62]. Metaheuristics follow the second path, generating and modifying candidate route sets with problem specific heuristic or evolutionary operators [38, 41]. More recently, learning based methods reduce manual design by learning policies that construct routes directly or heuristics that guide search [70, 24]. Across these approaches, local design choices remain hard to assess because their value depends on the completed route set.

This limitation motivates a route construction procedure that improves local decisions through lookahead over partial networks, while avoiding expensive simulation at every node. To address this, we introduce AlphaTransit, a search-guided framework that couples MCTS with a neural policy-value network, shown in Fig. 2. For each partial route, the policy provides priors over feasible extensions, and the value estimates downstream design quality. MCTS combines these predictions into a visit-count policy, from which the next extension is sampled to advance the design. During search, leaf states are evaluated by the value head rather than simulator rollouts. The simulator returns a terminal reward only after a complete route set is constructed; the resulting samples are used to update the network to match search-improved policies and predict final rewards. Our contributions are:

- We introduce AlphaTransit and show that learned lookahead is effective for simulator-defined bus TRNDP, achieving the highest service rate among evaluated methods in both demand regimes.
- We present a Bloomington TRNDP benchmark with a topologically accurate road network, census-derived origin-destination demand, and an existing human-designed bus transit reference.
- We isolate the role of learned lookahead by comparing AlphaTransit against End-to-End Reinforcement Learning (RL) without decision-time search and Pure MCTS without learned priors or value estimates, alongside heuristic, metaheuristic, neural-evolutionary, and real-world baselines.

We evaluate AlphaTransit on the Bloomington benchmark under mixed and full transit demand regimes. AlphaTransit reaches mean service rates of 54.64% and 82.08%, respectively, yielding relative service-rate gains of 9.9% and 11.4% over End-to-End RL, and 2.5% and 11.2% over Pure MCTS. Under mixed-demand, AlphaTransit also attains the highest bus utilization and second-highest route efficiency. Under full transit demand, it obtains the lowest wait time, highest route efficiency, and highest bus utilization. Together, these comparisons indicate that coupling learned priors with lookahead search makes sparse terminal feedback more useful for route design.

## 2 Related Work

TRNDP has been studied for over four decades, spanning exact, heuristic, and learning-based methods [28, 16, 43]. Exact and decomposition-based formulations provide optimization structure for analytical variants of the problem [5, 62]. Their tractability, however, depends on simplified demand, assignment, and operating models, which makes simulator-defined objectives with congestion and vehicle capacity difficult to optimize directly. Classical metaheuristics, including genetic algorithms [38], simulated annealing [71], and bee colony optimization [41, 16], handle larger design spaces and avoid some of these abstractions, but depend on tailored move operators, penalty terms, and instance-specific tuning. More recently, learning-based approaches reduce manual design by constructing routes with reinforcement learning [70, 15, 32] or by learning heuristics to guide evolutionary search [22, 23, 24]. These approaches move toward reusable policies, but most embed learning inside a larger metaheuristic loop rather than producing a standalone construction policy. For methods that construct routes directly, credit assignment remains difficult because the effect of a local route extension on network-wide passenger flow is only observed once routes have been assembled.

Most studies test on small or synthetic benchmark networks [34, 21] that do not reflect the scale or coupling of realistic urban networks, or rely on analytical objectives and assignment approximations. A smaller line of work uses simulation-based evaluation for bus-oriented transit network redesign and optimization [36, 42]. Both matter because congestion, vehicle capacity, transfer delays, and passenger reassignment are the mechanisms that make TRNDP rewards stochastic, delayed, and nonlocal. Such rewards favor methods with explicit lookahead, where search statistics both guide action selection and serve as training targets for a learned policy [52, 47]. Consistent with this motivation, MCTS has been applied in adjacent settings [27], including Pareto-optimal transit route planning [66], spatial network augmentation [14], surrogate-assisted combinatorial optimization [2], and metro network expansion combining deep RL with MCTS. The closest neural-search method is MetroZero [1], but it selects metro expansion stations under budget constraints, whereas AlphaTransit constructs complete bus route sets on existing shared roads and evaluates them through passenger and traffic simulation. However, to the best of our knowledge, no prior work on bus TRNDP brings these three strands together. Learned policy and value priors guide search, MCTS-based lookahead refines action selection, and simulation-based evaluation supplies the terminal reward signal. AlphaTransit integrates all three by training policy and value networks against targets produced by MCTS, with rewards drawn from mesoscopic traffic simulation on a city-scale road network [50].

## 3 AlphaTransit

### 3.1 Problem Formulation

We model the road network as an undirected graph  $G = (V, E)$ , where  $V = \{v_1, \dots, v_n\}$  denotes nodes such as intersections and  $E$  denotes bidirectional road segments, each with length  $\ell_e > 0$  and free-flow speed  $c_e > 0$ . Travel demand is exogenous and is represented by an origin–destination matrix  $D \in \mathbb{R}_{\geq 0}^{n \times n}$ , where  $D_{ij}$  gives the number of trips per hour from origin node  $i$  to destination

node  $j$ . The modal-split parameter  $\alpha \in [0, 1]$  assigns a fraction of total OD demand to bus transit, so that  $D_{ij}^{\text{tr}} = \alpha D_{ij}$ .

A transit route  $r_k = (v_{k,1}, \dots, v_{k,L_k})$  is an ordered simple path on  $G$ : consecutive nodes must share an edge, i.e.  $\{v_{k,q}, v_{k,q+1}\} \in E$ , and no node repeats within a route. Routes are operated bidirectionally, so buses traverse both  $v_{k,1} \rightarrow \dots \rightarrow v_{k,L_k}$  and its reverse. For each completed route set, bus passenger requests are evaluated on the induced transit graph. Passenger itineraries follow shortest-distance paths on this graph and may use one route or transfer across multiple routes to complete their journey. A complete transit design specifies an indexed  $K$ -tuple of routes,  $\Pi = (r_1, \dots, r_K)$ , together with stop spacing and service frequencies:

$$(\Pi, \mathbf{g}, \mathcal{F}) = ((r_1, \dots, r_K), (g_1, \dots, g_K), (\mathcal{F}_1, \dots, \mathcal{F}_K)), \quad (1)$$

where  $g_k \geq 1$  defines the stop spacing for route  $k$ , and  $\mathcal{F}_k \in \mathbb{N}_{>0}$  is the service frequency in buses per hour. The tuple notation provides the index-wise correspondence between route  $r_k$ , spacing  $g_k$ , and frequency  $\mathcal{F}_k$ ; the simulator reward is invariant to permutations of route order. Given  $G$ ,  $D$ , and design constraints such as the route count  $K$  and route length bounds, the Transit Route Network Design Problem seeks to maximize expected system-level performance:

$$(\Pi^*, \mathbf{g}^*, \mathcal{F}^*) \in \operatorname{argmax}_{(\Pi, \mathbf{g}, \mathcal{F}) \in \mathcal{X}} \mathbb{E}_\xi [\mathcal{R}(\Pi, \mathbf{g}, \mathcal{F}; \xi)], \quad (2)$$

where  $\mathcal{X}$  is the feasible design set,  $\xi$  denotes simulator stochasticity, and  $\mathcal{R}(\cdot; \xi)$  is a traffic-simulation performance measure. We fix the stop spacing to  $g_k = 1$  for all routes and assign frequencies using a max-load rule (Appendix B); thus, the learned decision variables reduce to the route tuple  $\Pi$ .

We cast the route construction as a finite-horizon Markov decision process  $(\mathcal{S}, \mathcal{A}, P, R_T, s_0)$ , where  $\mathcal{S}$  is the state space,  $\mathcal{A} = \{1, \dots, n\}$  is the action space with infeasible actions masked at each step,  $P$  is the transition function, and  $s_0$  is the initial state. Since stop spacing and frequencies are deterministic once  $\Pi$  is chosen, the terminal reward is given by

$$R_T(\Pi; \xi) = \mathcal{R}(\Pi, \mathbf{1}, \mathcal{F}(\Pi); \xi), \quad (3)$$

Here  $\mathcal{F}(\Pi) = F^{\text{ml}}(\Pi)$  is the max-load frequency projection defined in Appendix B. For a completed route set, this projection is the componentwise minimal positive frequency vector satisfying fixed-load capacity constraints. Appendix B proves this projection property and defines the corresponding gap between the projected route-only objective and joint route-frequency optimization. Let  $\mu_\theta$  be the executed construction policy. For End-to-End RL,  $\mu_\theta(a | s) = \pi_\theta(a | s)$ ; for AlphaTransit,  $\mu_\theta$  is the MCTS policy induced by the policy-value network  $f_\theta$ . The learning objective is

$$J(\theta) = \mathbb{E}_{\Pi \sim \mu_\theta, \xi} [R_T(\Pi; \xi)], \quad (4)$$

where  $\Pi$  is the complete route set constructed under  $\mu_\theta$ . In each episode, the agent sequentially constructs  $K$  routes, each satisfying  $|r_k| \leq L_{\text{max}}$ , by extending one route at a time. In our setting,  $K = 16$  and  $L_{\text{max}} = 14$ . Routes start at a transit-center hub and grow one node at a time from the frontier, the most recently appended node, with feasible actions restricted to one-hop neighbors not yet visited in the current route. The policy has no separate stop action; a route is finalized automatically when it reaches  $L_{\text{max}}$  or when no valid extension exists. Thus, the learned decision horizon is bounded by  $T_{\text{max}} = K(L_{\text{max}} - 1)$  and may be shorter when routes terminate early.

**State.** At each step  $t$ , the state encodes the road graph and the transit network constructed so far, including completed routes, the current partial route, and the frontier node. The policy network receives this as  $s_t = (X_t, \mathcal{I}, Z)$ , where  $\mathcal{I}$  is a directed edge list,  $Z$  stores edge features such as length and free-flow speed, and  $X_t \in \mathbb{R}^{n \times d_x}$  contains per-node features capturing spatial attributes, OD demand aggregates, route membership, and frontier status. The full specification is in Appendix C. **Action.** At each step the agent selects  $a_t \in \mathcal{A} = \{1, \dots, n\}$ , with  $a_t = i$  appending  $v_i$  to the current route. Let  $v_{f_t}$  be the frontier node of partial route  $r_{k,t}$ . The admissible candidate set is

$$\mathcal{C}_t = \{i \in \{1, \dots, n\} : \{v_{f_t}, v_i\} \in E \text{ and } v_i \notin r_{k,t}\}, \quad (5)$$

so every valid action extends the route by one hop while preserving the simple-path constraint, and  $\mathcal{C}_t = \emptyset$  triggers early termination. Because  $\mathcal{C}_t$  is restricted to unvisited one-hop neighbors of the frontier and road graphs are sparse, typically  $|\mathcal{C}_t| \ll n$ . We therefore apply invalid-action

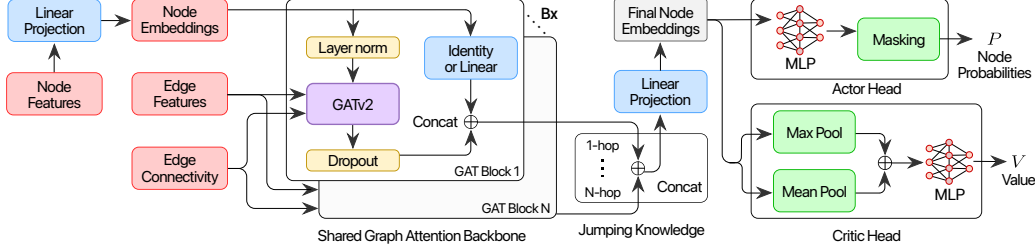


Figure 3: AlphaTransit policy-value network. Node features are projected to embeddings, then passed through a GATv2 backbone using edge connectivity and edge attributes. Successive attention blocks produce multi-hop representations, which Jumping Knowledge aggregation concatenates and projects to node embeddings. A node-wise actor MLP produces logits, from which infeasible actions are masked. A graph-level critic uses global mean and max pooling and predicts  $V_\theta(s)$ . AlphaTransit uses  $P_\theta(\cdot | s)$  as MCTS action priors and  $V_\theta(s)$  as the leaf-value estimate.

masking [26]: a binary mask  $m_t \in \{0, 1\}^n$  with  $m_t[i] = 1\{i \in \mathcal{C}_t\}$  restricts the policy to feasible extensions. Given logits  $\ell_\theta(s_t) \in \mathbb{R}^n$ , the masked policy, for  $\mathcal{C}_t \neq \emptyset$ , is

$$\pi_\theta(a_t=i | s_t) = \begin{cases} \frac{\exp(\ell_i)}{\sum_{j: m_t[j]=1} \exp(\ell_j)} & \text{if } m_t[i] = 1, \\ 0 & \text{if } m_t[i] = 0. \end{cases} \quad (6)$$

Node feasibility is also exposed in  $X_t$  via a valid-next flag, allowing the policy to observe which actions are admissible while the mask enforces zero probability on invalid actions.

**Reward.** For a complete route design  $\Pi = (r_1, \dots, r_K)$ , we assign frequencies  $\mathcal{F}(\Pi)$  via the max-load rule and evaluate the design through simulation. The terminal reward for training is

$$R_T(\Pi; \xi) = b_0 \Psi + b_1 \rho - b_2 \hat{t}_{\text{wait}} - b_3 \hat{t}_{\text{move}} - b_4 \omega - b_5 \frac{N_{\text{bus}}}{K} + b_6 u, \quad (7)$$

with  $b_0=60$ ,  $b_1=45$ ,  $b_2=20$ ,  $b_3=10$ ,  $b_4=10$ ,  $b_5=2$ ,  $b_6=12$ , chosen to balance passenger and operator objectives; planners with different goals can set different weights. The coverage potential  $\Psi$  is the fraction of total OD demand assigned to transit and reachable under  $\Pi$ . Let  $G_\Pi = (V_\Pi, E_\Pi)$  be the graph induced by  $\Pi$ , with  $V_\Pi$  as served nodes and  $E_\Pi$  as edges. Then

$$\Psi = \frac{\sum_{(i,j) \in \mathcal{P}_{\text{reach}}} D_{ij}^{\text{tr}}}{\sum_{i,j \in V} D_{ij}}, \quad \mathcal{P}_{\text{reach}} = \{(i, j) : i, j \in V_\Pi \text{ and } j \text{ is reachable from } i \text{ in } G_\Pi\}. \quad (8)$$

The service term is  $\rho = N_{\text{boarded}}/N_{\text{OD}}$ , where  $N_{\text{boarded}}$  counts passengers who complete trips or are onboard at the simulation horizon and  $N_{\text{OD}}$  is citywide OD demand over that horizon. This fixed-denominator reward term differs from the reported service rate  $\sigma = N_{\text{boarded}}/N_{\text{want}}$  in Appendix G. Let  $\bar{t}_{\text{wait}}$  and  $\bar{t}_{\text{move}}$  denote raw average waiting and in-vehicle times in minutes over served passengers. The reward uses capped normalized time penalties,

$$\hat{t}_{\text{wait}} = \min\{\bar{t}_{\text{wait}}/30, 1\}, \quad \hat{t}_{\text{move}} = \min\{\bar{t}_{\text{move}}/40, 1\}.$$

The remaining terms are route overlap ratio  $\omega$ , fleet size  $N_{\text{bus}}$ , and bus utilization  $u$ . The scalar terminal reward is separately normalized online when forming the value target [25, 12].

### 3.2 Search-Guided Reinforcement Learning

Sequential transit design provides a sparse training signal because node extensions are evaluated only after route designs are completed, frequencies are assigned, and passenger flows are simulated. Local construction policies, including heuristics and end-to-end RL, select each node without looking ahead to downstream network effects at decision time; these effects are observed only after design evaluation. AlphaTransit runs Monte Carlo Tree Search (MCTS) at each construction state, guided by the policy-value network  $f_\theta$ . The visit counts define an MCTS policy used both to sample the next node and to train the actor. This makes AlphaTransit a search-guided reinforcement learning method.

Fig. 3 shows  $f_\theta$  mapping graph state  $s = (X, \mathcal{I}, Z)$  to  $f_\theta(s) = (P_\theta(\cdot | s), V_\theta(s))$ , where  $P_\theta$  gives node-action probabilities and  $V_\theta$  estimates the value of completing a design from  $s$ . The network

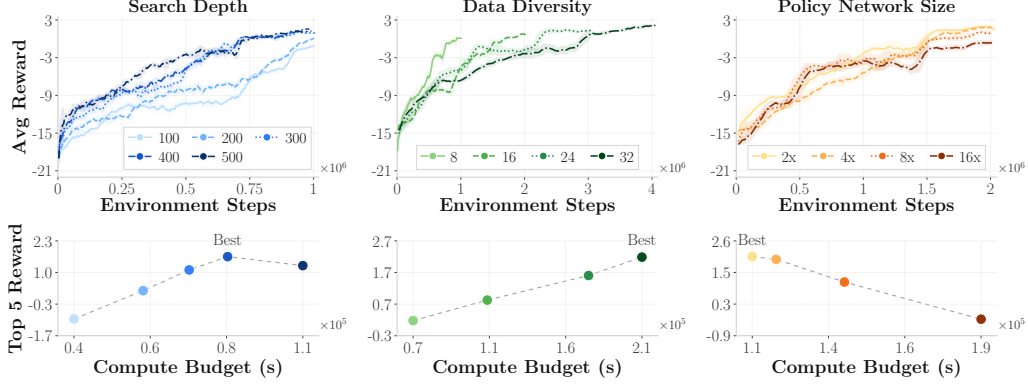


Figure 4: AlphaTransit scaling behavior under mixed demand ( $\alpha = 0.3$ ). Search depth, data diversity, and policy size show different quality–compute trade-offs. **LEFT:** Performance peaks at  $N_{\text{iter}} = 400$  with top-5 reward 1.67, rather than increasing monotonically, while  $N_{\text{iter}} = 500$  lowers reward to 1.29 with higher runtime. **MIDDLE:** Increasing episodes per iteration gives the highest reward, 2.20 at 32 episodes, but costs the most compute. **RIGHT:** Policy-size labels ( $2\times$ - $16\times$ ) denote the GAT-block repetition count; the  $2\times$  policy reaches 2.05, while the  $16\times$  policy falls below 0.

projects node features and processes them with a shared GATv2 [7] backbone using  $\mathcal{I}$  and  $\mathcal{Z}$ . The actor and critic share this backbone. For fixed hidden widths, attention heads, and block count  $B$ , the parameter count is independent of  $n$ , while a forward pass costs  $O(B(n + |\mathcal{I}|))$ , linear in  $n$  for sparse road graphs. Jumping Knowledge aggregation [68] gives both heads access to multiple receptive-field scales. Let  $\mathbf{h}_v^{(b)}$  denote the representation of node  $v$  after block  $b$ . The block outputs are concatenated and projected to a node embedding:

$$\mathbf{z}_v = W_{\text{JK}} \left[ \mathbf{h}_v^{(1)} \parallel \dots \parallel \mathbf{h}_v^{(B)} \right] + \mathbf{b}_{\text{JK}}, \quad (9)$$

where  $W_{\text{JK}} \in \mathbb{R}^{d \times \sum_{b=1}^B d_b}$  maps the concatenation to dimension  $d$ . The actor applies the same MLP to each  $\mathbf{z}_v$ , producing one logit per node, then masks infeasible logits via Eq. 6. The critic concatenates global mean and max pooled node embeddings and maps the pooled representation to  $V_\theta(s)$ . Appendix D provides architecture details.

At each construction state  $s_t$ , AlphaTransit runs an MCTS rooted at  $s_t$  over admissible route-extension actions. For any tree state  $s$ , let  $\mathcal{C}(s)$  be the candidate set from Eq. 5; at the root,  $\mathcal{C}(s_t) = \mathcal{C}_t$ . Each edge  $(s, a)$ , with  $a \in \mathcal{C}(s)$ , stores  $N(s, a)$ ,  $W(s, a)$ ,  $Q(s, a)$ , and prior  $P_\theta(a | s)$ , where  $Q(s, a) = W(s, a)/N(s, a)$  if  $N(s, a) > 0$  and  $Q(s, a) = 0$  otherwise. Leaf states are evaluated by  $V_\theta$ , avoiding simulator calls inside the tree. Each simulation then proceeds through:

- **Selection:** Starting at the root, each simulation selects actions using the PUCT rule [45], which combines the current value estimate with an exploration bonus:

$$a^* = \operatorname{argmax}_{a \in \mathcal{C}(s)} \left[ Q(s, a) + c P_\theta(a | s) \frac{\sqrt{1 + \sum_{b \in \mathcal{C}(s)} N(s, b)}}{1 + N(s, a)} \right]. \quad (10)$$

Here,  $c$  balances exploitation of high-value actions with exploration of less-visited ones.

- **Expansion and evaluation:** When selection reaches an unexpanded leaf  $s_{\text{leaf}}$ , a single forward pass of  $f_\theta$  produces priors  $P_\theta(a | s_{\text{leaf}})$  for each  $a \in \mathcal{C}(s_{\text{leaf}})$  and a scalar value  $V_\theta(s_{\text{leaf}})$  used as the backed-up leaf value. No simulator rollout is performed inside the tree; the simulator is invoked only after the executed trajectory completes a full route set.
- **Backpropagation:** After evaluation, statistics are updated along the path from leaf to root:  $N(s, a) \leftarrow N(s, a) + 1$ ,  $W(s, a) \leftarrow W(s, a) + V_\theta(s_{\text{leaf}})$ , and  $Q(s, a) \leftarrow W(s, a)/N(s, a)$ .

After  $N_{\text{iter}}$  simulations, the action distribution at the root is derived from visit counts:

$$\pi_t(a | s_t) = \frac{N(s_t, a)^{1/\tau}}{\sum_{b \in \mathcal{C}(s_t)} N(s_t, b)^{1/\tau}}, \quad a \in \mathcal{C}(s_t), \quad (11)$$



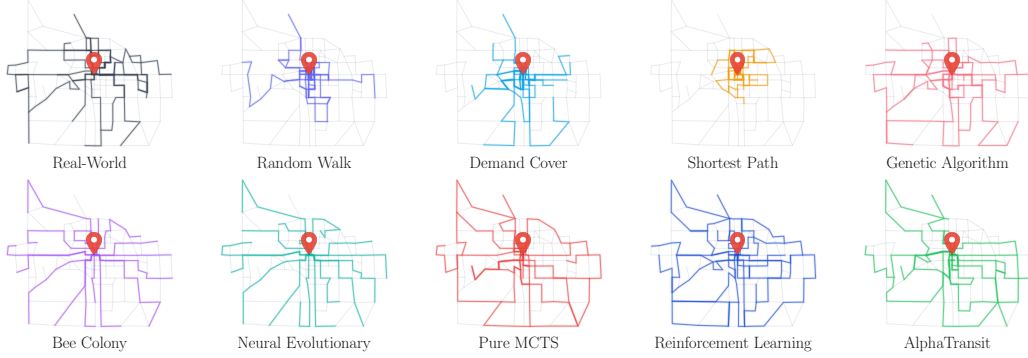


Figure 6: Selected route designs under mixed demand ( $\alpha = 0.3$ ) on the Bloomington network, with the red pin marking the transit center. Each panel overlays a method’s routes on the same street basemap. AlphaTransit covers 117 nodes (81.8%), has 24.4% shared-edge overlap, and spans 120.8 km. Real-World covers 114 nodes (79.7%), has 34.8% overlap, and spans 115.6 km, while End-to-End RL covers 138 nodes (96.5%), has 19.0% overlap, and spans 142.3 km. AlphaTransit is therefore more targeted than End-to-End RL.

served demand to bus transit, and  $\alpha = 0.3$ , which reflects a typical urban public-transport share [9]. We train one AlphaTransit policy per scenario. All completed route sets are evaluated through the same UXsim reporting pipeline, with frequencies assigned by the same max-load rule and metrics derived from Eq. 7. Training, hyperparameters, and baseline details are in Appendix E. We compare against the following baselines:

- **Real-World:** The 16 bus routes operated by Bloomington Transit [6]. The agency’s design reflects broader objectives such as equity, coverage, and budget that extend beyond our simulator-defined criteria; we include it as a real-world reference.
- **Random Walk:** Samples next node uniformly from the admissible candidate set  $\mathcal{C}_t$ .
- **Demand Coverage:** Samples next node  $i \in \mathcal{C}_t$  proportional to its demand interaction with the partial route,  $s(i) = \sum_{j \in r_k} (D_{ij} + D_{ji})$ .
- **Shortest Path:** Samples next node  $i \in \mathcal{C}_t$  inversely proportional to edge length,  $s(i) = 1/\ell_{u,i}$ , where  $u$  is the frontier and  $\ell_{u,i}$  is the length of edge  $\{u, i\}$ .
- **Genetic Algorithm:** A widely used metaheuristic for transit network design [18, 39]. Each individual represents a complete route tuple  $\Pi = (r_1, \dots, r_K)$  that evolves over generations through tournament selection, route-exchange crossover, and path-regeneration mutation.
- **Bee Colony Optimization:** A swarm-intelligence metaheuristic [41] that evolves a population of route sets through two mutations: (i) demand-weighted shortest-path route replacement and (ii) node addition/deletion at route endpoints. We adapt the implementation from [24], preserve the analytical objective during route generation, and evaluate the completed routes through UXsim.
- **Neural Evolutionary Algorithm:** Augments Bee Colony Optimization with mutations proposed by a graph neural policy trained via RL [23, 24]. We train the policy under our setup, preserving the original analytical objective during search, and evaluate the completed routes through UXsim.
- **Pure Monte Carlo Tree Search [13, 8]:** Combines uniform action priors  $P(s, a) = 1/|\mathcal{C}_t|$  and full UXsim rollouts. The PUCT formula (Eq. 10) and search budget match AlphaTransit.
- **End-to-End RL:** A direct policy-learning baseline trained with PPO [49] and GAE [48]. It selects actions from the masked policy  $\pi_\theta(a | s_t)$  without decision-time search. It uses the same state, action mask, terminal reward, and policy architecture as AlphaTransit, plus lightweight shaping rewards with  $\gamma = 0.999$  to make long horizon policy gradient training feasible.

We evaluate network designs across three axes from transit evaluation criteria [64, 65]. The first weighs passenger benefit against operator cost: service rate (% of potential demand boarded) vs. fleet size. Here, potential demand refers to the reachable fixed-transit-demand denominator  $N_{\text{want}}$  defined in Appendix G. The second groups two passenger-experience metrics: total journey time (minutes), elapsed passenger time including waiting and in-vehicle movement, and transfer rate (% of trips requiring a transfer). The third groups two operator-side efficiency metrics: route efficiency (passengers served per km of route) and bus utilization (% of bus occupied). Full definitions are provided in Appendix G.

Method	Passenger Metrics				Operator Metrics			
	Service Rate (%) $\uparrow$	Wait Time (min) $\downarrow$	Transfer Rate (%) $\downarrow$	Journey Time (min) $\downarrow$	Route Efficiency $\uparrow$	Fleet Size $\downarrow$	Bus Util. (%) $\uparrow$	
$\alpha = 0.3$	Demand Cover	73.80 $\pm$ 5.40	15.93 $\pm$ 2.89	49.73 $\pm$ 4.64	41.03 $\pm$ 3.71	103.30 $\pm$ 18.39	96.10 $\pm$ 16.23	38.03 $\pm$ 3.93
	Shortest Path	90.72 $\pm$ 7.53	9.91 $\pm$ 3.12	57.48 $\pm$ 4.03	29.62 $\pm$ 2.80	129.85 $\pm$ 37.78	72.20 $\pm$ 21.67	23.77 $\pm$ 3.63
	End-to-End RL	88.96 $\pm$ 2.01	7.38 $\pm$ 0.71	58.91 $\pm$ 1.12	39.99 $\pm$ 1.21	225.69 $\pm$ 9.93	277.00 $\pm$ 0.00	47.34 $\pm$ 3.14
	AlphaTransit	89.25 $\pm$ 0.84	7.07 $\pm$ 0.54	57.70 $\pm$ 0.70	38.19 $\pm$ 1.00	200.87 $\pm$ 4.77	241.00 $\pm$ 0.00	42.14 $\pm$ 1.18
$\alpha = 1.0$	Demand Cover	62.00 $\pm$ 7.96	12.41 $\pm$ 4.65	41.28 $\pm$ 5.80	45.37 $\pm$ 6.01	256.48 $\pm$ 45.16	270.90 $\pm$ 32.43	45.87 $\pm$ 4.34
	Shortest Path	79.69 $\pm$ 16.75	7.21 $\pm$ 3.17	49.38 $\pm$ 10.86	33.08 $\pm$ 6.32	324.53 $\pm$ 101.56	190.60 $\pm$ 55.88	30.55 $\pm$ 5.37
	End-to-End RL	55.03 $\pm$ 1.79	15.80 $\pm$ 0.74	46.76 $\pm$ 0.97	50.68 $\pm$ 1.25	313.00 $\pm$ 7.73	633.00 $\pm$ 0.00	54.79 $\pm$ 1.07
	AlphaTransit	90.72 $\pm$ 0.71	6.03 $\pm$ 0.58	63.33 $\pm$ 0.48	35.50 $\pm$ 0.61	396.30 $\pm$ 6.27	330.00 $\pm$ 0.00	44.34 $\pm$ 0.32

Table 1: Cross-city transfer on Laval. Values report mean  $\pm$  standard deviation over 10 seeds; Laval-optimized learning/search methods are omitted to isolate fixed-policy transfer. Under full transit demand ( $\alpha = 1.0$ ), AlphaTransit reaches 90.72% service rate, versus 55.03% for End-to-End RL, and also obtains the lowest wait time and highest route efficiency among transfer-compatible methods, giving the clearest advantage. Under mixed demand ( $\alpha = 0.3$ ), AlphaTransit remains close to Shortest Path in service rate while giving the lowest wait time.

**Results.** Fig. 1 summarizes the learning dynamics and decision-time search cost under mixed demand. End-to-End RL relies on dense shaping to learn over the long route-construction horizon: the Delta-coverage variants outperform terminal-only and Raw + Early Stopping feedback, with Delta + No Early Stopping used in later comparisons. AlphaTransit improves sample efficiency under the same environment-step budget because MCTS provides search-improved action targets, while learned policy-value estimates keep decision-time search in seconds rather than the hundreds to thousands of seconds required by Pure MCTS. Fig. 4 focuses on AlphaTransit scaling: more compute helps only when allocated carefully. Performance improves with search depth up to  $N_{\text{iter}} = 400$  and drops at 500; episodes per iteration raise top-5 reward from 0.19 to 2.20; larger GAT policies provide no gain.

Fig. 5 gives the mixed-demand ( $\alpha = 0.3$ ) comparison on the Bloomington benchmark, and Appendix G reports the complete metric table and complementary analyses. Under mixed demand, AlphaTransit achieves the highest service rate, 54.64%, with 80 buses and the highest bus utilization, 22.10%. The same service-rate ranking holds under full transit demand ( $\alpha = 1.0$ ), where AlphaTransit reaches 82.08% service rate and also obtains the best wait time, route efficiency, and bus utilization. Since these metrics capture different passenger and operator objectives, no method dominates every axis: AlphaTransit is strongest on service rate and utilization, while other methods can obtain lower transfer rates, shorter journey times, or smaller fleets in some settings.

To test whether the gain comes from combining learning with search, we compare against End-to-End RL, which removes decision-time search, and Pure MCTS, which uses the same search budget without learned prior-value estimates. Relative to End-to-End RL, AlphaTransit improves service rate by 9.9% at  $\alpha = 0.3$  and 11.4% at  $\alpha = 1.0$ ; relative to Pure MCTS, it improves service rate by 2.5% and 11.2%, respectively. At  $\alpha = 0.3$ , Pure MCTS reaches 53.30% with 86 buses, whereas AlphaTransit reaches 54.64% with 80 buses, suggesting search works best with learned estimates. Fig. 6 shows that the gain is not simply broader coverage: AlphaTransit covers fewer nodes and less route distance than End-to-End RL (117 vs. 138 nodes; 120.8 km vs. 142.3 km), yet serves more demand with a smaller fleet. Compared with Real-World routes, it has similar node coverage but reduces shared edge overlap (24.4% vs. 34.8%) and improves service rate from 42.77% to 54.64%.

Finally, Table 1 evaluates cross-city transfer on the larger Laval network using policies trained only on Bloomington. Under full transit demand ( $\alpha = 1.0$ ), AlphaTransit reaches 90.72% service rate, compared with 55.03% for End-to-End RL, and also obtains the lowest wait time and highest route efficiency among the listed transfer-compatible methods. Under mixed demand ( $\alpha = 0.3$ ), Shortest Path has the highest service rate, while AlphaTransit remains close at 89.25% and gives the lowest wait time. The Laval results are mixed across metrics, but they show that the search-guided policy transfers to a larger network, with the clearest advantage in the higher-demand setting.

## 5 Conclusion and Discussion

Transit route network design is a sequential combinatorial problem in which each route-extension decision is judged only after the full network is assembled and simulated. This delayed and nonlocal

feedback makes TRNDP deceptive, since extensions that look useful locally can later create redundant overlap, transfer burden, or capacity bottlenecks. AlphaTransit addresses this challenge by coupling Monte Carlo Tree Search with a graph attention policy-value network. The policy proposes feasible route extensions, the value estimates downstream design quality, and search refines each decision without simulator rollouts inside the tree.

We also introduce a new Bloomington benchmark with realistic road topology, census-derived origin-destination demand, and existing transit routes. The results show that AlphaTransit achieves the highest service rate in both mixed and full transit demand, reaching 54.64% at  $\alpha = 0.3$  and 82.08% at  $\alpha = 1.0$ . Relative to End-to-End Reinforcement Learning, AlphaTransit improves service rate by 9.9% and 11.4%; relative to Pure Monte Carlo Tree Search, it improves service rate by 2.5% and 11.2%. The scaling results further show that calibrated search depth and episode diversity matter more than simply enlarging the policy network. Together, these results establish learned lookahead as an effective mechanism for coordinating route construction under delayed, network-level evaluation.

Several limitations of the current work exist. First, the geographic scope is limited: most experiments and model development rely on the Bloomington benchmark, so the results may not fully capture the diversity of transit networks across cities. The Laval experiment provides an initial cross-city check, but broader validation on additional metropolitan systems is needed. A second limitation is the route-start assumption. Every route begins at the transit-center hub, which matches Bloomington Transit and reduces the construction space, but may not fit multi-hub, crosstown, or grid-like networks. Finally, the simulator and decision model abstract several deployment factors: demand is represented by a static peak-hour OD matrix, and the reward does not explicitly encode equity, accessibility, reliability, budget, or robustness constraints.

Future work should extend training and evaluation with endogenous mode choice and to more cities and larger metropolitan networks, relax the transit-center start constraint, or learn route origins jointly with route extensions. Other promising extensions include time-varying demand, stochastic disruptions, a second-stage frequency policy, and constrained objectives that more clearly expose service-quality trade-offs across neighborhoods.

## References

- [1] Khalid Alkilane and Der-Horng Lee. MetroZero: Deep reinforcement learning and Monte Carlo tree search for optimized metro network expansion. *IEEE Transactions on Intelligent Transportation Systems*, 26(1):810–823, 2025.
- [2] Saeid Amiri, Parisa Zehtabi, Danial Dervovic, and Michael Cashmore. Surrogate assisted monte carlo tree search in combinatorial optimization. *arXiv preprint arXiv:2403.09925*, 2024.
- [3] American Public Transportation Association. Public transportation ridership update. Policy brief, APTA, 2025.
- [4] Irwan Bello, Hieu Pham, Quoc V Le, Mohammad Norouzi, and Samy Bengio. Neural combinatorial optimization with reinforcement learning. *arXiv preprint arXiv:1611.09940*, 2016.
- [5] Dimitris Bertsimas, Yee Sian Ng, and Julia Yan. Data-driven transit network design at scale. *Operations Research*, 69(4):1118–1133, 2021.
- [6] Bloomington Transit. Gtfs schedule dataset. <https://bloomingtontransit.com/gtfs/>.
- [7] Shaked Brody, Uri Alon, and Eran Yahav. How attentive are graph attention networks? *arXiv preprint arXiv:2105.14491*, 2021.
- [8] Cameron B Browne, Edward Powley, Daniel Whitehouse, Simon M Lucas, Peter I Cowling, Philipp Rohlfshagen, Stephen Tavener, Diego Perez, Spyridon Samothrakis, and Simon Colton. A survey of monte carlo tree search methods. *IEEE Transactions on Computational Intelligence and AI in games*, 4(1):1–43, 2012.
- [9] Ralph Buehler, John Pucher, and Oliver Dümmler. Verkehrsverbund: The evolution and spread of fully integrated regional public transport in germany, austria, and switzerland. *International Journal of Sustainable Transportation*, 13(1):36–50, 2019.
- [10] Cambridge Systematics, Inc. and Texas Transportation Institute. Traffic congestion and reliability: Trends and advanced strategies for congestion mitigation. Final Report FHWA-HOP-05-064, Federal Highway Administration, Washington, DC, September 2005.
- [11] Avishai Ceder. *Public transit planning and operation: Modeling, practice and behavior*. CRC press, 2016.
- [12] Jeremiah Coholich. A bag of tricks for deep reinforcement learning, 2023.
- [13] Rémi Coulom. Efficient selectivity and backup operators in monte-carlo tree search. In *International conference on computers and games*, pages 72–83. Springer, 2006.
- [14] Victor-Alexandru Darvari, Stephen Hailes, and Mirco Musolesi. Planning spatial networks with Monte Carlo tree search. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 479(2269):20220383, 2023.
- [15] Ahmed Darwish, Momen Khalil, and Karim Badawi. Optimising public bus transit networks using deep reinforcement learning. In *2020 IEEE 23rd International Conference on Intelligent Transportation Systems (ITSC)*, pages 1–7. IEEE, 2020.
- [16] Javier Durán-Micco and Pieter Vansteenwegen. A survey on the transit network design and frequency setting problem. *Public transport*, 14(1):155–190, 2022.
- [17] Wei Fan. *Optimal transit route network design problem: Algorithms, implementations, and numerical results*. The University of Texas at Austin, 2004.
- [18] Wei Fan and Randy B Machemehl. Optimal transit route network design problem with variable transit demand: genetic algorithm approach. *Journal of transportation engineering*, 132(1):40–51, 2006.
- [19] Alhussein Fawzi, Matej Balog, Aja Huang, Thomas Hubert, Bernardino Romera-Paredes, Mohammadamin Barekatin, Alexander Novikov, Francisco J R. Ruiz, Julian Schrittwieser, Grzegorz Swirszcz, et al. Discovering faster matrix multiplication algorithms with reinforcement learning. *Nature*, 610(7930):47–53, 2022.

- [20] Peter G Furth and Nigel HM Wilson. Setting frequencies on bus routes: Theory and practice. *Transportation Research Record*, 818(1981):1–7, 1981.
- [21] Philipp Heyken Soares, Christine L Mumford, Kwabena Amponsah, and Yong Mao. An adaptive scaled network for public transport route optimisation. *Public Transport*, 11(2):379–412, 2019.
- [22] Andrew Holliday and Gregory Dudek. Augmenting transit network design algorithms with deep learning. In *2023 IEEE 26th International Conference on Intelligent Transportation Systems (ITSC)*, pages 2343–2350. IEEE, 2023.
- [23] Andrew Holliday and Gregory Dudek. A neural-evolutionary algorithm for autonomous transit network design. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, pages 4457–4464. IEEE, 2024.
- [24] Andrew Holliday, Ahmed El-Geneidy, and Gregory Dudek. Learning heuristics for transit network design and improvement with deep reinforcement learning. *Transportmetrica B: Transport Dynamics*, 13(1):2561863, 2025.
- [25] Shengyi Huang, Rousslan Fernand Julien Dossa, Antonin Raffin, Anssi Kanervisto, and Weixun Wang. The 37 implementation details of proximal policy optimization. *The ICLR Blog Track 2023*, 2022.
- [26] Shengyi Huang and Santiago Ontañón. A closer look at invalid action masking in policy gradient algorithms. *arXiv preprint arXiv:2006.14171*, 2020.
- [27] Marco Kemmerling, Daniel Lütticke, and Robert H Schmitt. Beyond games: a systematic review of neural monte carlo tree search applications. *arXiv preprint arXiv:2303.08060*, 2023.
- [28] Konstantinos Kepaptsoglou and Matthew Karlaftis. Transit route network design problem. *Journal of transportation engineering*, 135(8):491–505, 2009.
- [29] Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In *International Conference on Learning Representations*, 2019.
- [30] Yeong-Dae Kwon, Jinho Choo, Byoungjip Kim, Iljoo Yoon, Youngjune Gwon, and Seungjai Min. Pomo: Policy optimization with multiple optima for reinforcement learning. *Advances in neural information processing systems*, 33:21188–21198, 2020.
- [31] Eugene L Lawler. The traveling salesman problem: a guided tour of combinatorial optimization. *Wiley-Interscience Series in Discrete Mathematics*, 1985.
- [32] Junjun Li, Hao Dong, Xuedong Zhao, Hao Tang, Aimin Yin, and Ruchen Xue. A transit network design and frequency setting model with graph neural network and deep reinforcement learning. In *Sixth International Conference on Computer Information Science and Application Technology (CISAT 2023)*, volume 12800, page 128005Y. SPIE, 2023.
- [33] Pablo Alvarez Lopez, Michael Behrisch, Laura Bieker-Walz, Jakob Erdmann, Yun-Pang Flötteröd, Robert Hilbrich, Leonhard Lücken, Johannes Rummel, Peter Wagner, and Evamarie Wießner. Microscopic traffic simulation using SUMO. In *2018 21st International Conference on Intelligent Transportation Systems (ITSC)*, pages 2575–2582. IEEE, November 2018.
- [34] Christoph E Mandl. Evaluation and optimization of urban public transportation networks. *European Journal of Operational Research*, 5(6):396–404, 1980.
- [35] Daniel J Mankowitz, Andrea Michi, Anton Zhernov, Marco Gelmi, Marco Selvi, Cosmin Paduraru, Edouard Leurent, Shariq Iqbal, Jean-Baptiste Lespiau, Alex Ahern, et al. Faster sorting algorithms discovered using deep reinforcement learning. *Nature*, 618(7964):257–263, 2023.
- [36] Patrick Manser, Henrik Becker, Sebastian Hörl, and Kay W Axhausen. Designing a large-scale public transport network using agent-based microsimulation. *Transportation Research Part A: Policy and Practice*, 137:1–15, 2020.

- [37] Nancy McGuckin, Anthony Fucci, et al. Summary of travel trends: 2017 national household travel survey. Technical report, United States. Department of Transportation. Federal Highway Administration, 2018.
- [38] Christine L Mumford. New heuristic and evolutionary operators for the multi-objective urban transit routing problem. In *2013 IEEE congress on evolutionary computation*, pages 939–946. IEEE, 2013.
- [39] Muhammad Ali Nayeem, Md Khaledur Rahman, and M Sohel Rahman. Transit network design by genetic algorithm with elitism. *Transportation Research Part C: Emerging Technologies*, 46:30–45, 2014.
- [40] Gordon Frank Newell. A simplified car-following theory: a lower order model. *Transportation Research Part B: Methodological*, 36(3):195–205, 2002.
- [41] Miloš Nikolić and Dušan Teodorović. Transit network design by bee colony optimization. *Expert Systems with Applications*, 40(15):5945–5955, 2013.
- [42] Obiora A Nnene, Johan W Joubert, and Mark HP Zuidgeest. A simulation-based optimization approach for designing transit networks. *Public Transport*, 15(2):377–409, 2023.
- [43] Mahmoud Owais. Transit network design problem: a half century of methodological research. *Innovative Infrastructure Solutions*, 11(1):3, 2026.
- [44] R.P. Roess, E.S. Prassas, and W.R. McShane. *Traffic Engineering*. Pearson, 2011.
- [45] Christopher D Rosin. Multi-armed bandits with episode context. *Annals of Mathematics and Artificial Intelligence*, 61(3):203–230, 2011.
- [46] Marie Schmidt and Anita Schöbel. Planning and optimizing transit lines. *arXiv preprint arXiv:2405.10074*, 2024.
- [47] Julian Schrittwieser, Ioannis Antonoglou, Thomas Hubert, Karen Simonyan, Laurent Sifre, Simon Schmitt, Arthur Guez, Edward Lockhart, Demis Hassabis, Thore Graepel, Timothy Lillicrap, and David Silver. Mastering Atari, Go, chess and shogi by planning with a learned model. *Nature*, 588(7839):604–609, 2020.
- [48] John Schulman, Philipp Moritz, Sergey Levine, Michael Jordan, and Pieter Abbeel. High-dimensional continuous control using generalized advantage estimation. *arXiv preprint arXiv:1506.02438*, 2015.
- [49] John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.
- [50] Toru Seo. Uxsim: lightweight mesoscopic traffic flow simulator in pure python. *Journal of Open Source Software*, 10(106):7617, 2025.
- [51] David Silver, Aja Huang, Chris J. Maddison, Arthur Guez, Laurent Sifre, George van den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, Sander Dieleman, Dominik Grewe, John Nham, Nal Kalchbrenner, Ilya Sutskever, Timothy Lillicrap, Madeleine Leach, Koray Kavukcuoglu, Thore Graepel, and Demis Hassabis. Mastering the game of Go with deep neural networks and tree search. *Nature*, 529:484–489, 2016.
- [52] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharsan Kumaran, Thore Graepel, Timothy Lillicrap, Karen Simonyan, and Demis Hassabis. A general reinforcement learning algorithm that masters chess, shogi, and Go through self-play. *Science*, 362(6419):1140–1144, 2018.
- [53] David Silver, Julian Schrittwieser, Karen Simonyan, Ioannis Antonoglou, Aja Huang, Arthur Guez, Thomas Hubert, Lucas Baker, Matthew Lai, Adrian Bolton, Yutian Chen, Timothy Lillicrap, Fan Hui, Laurent Sifre, George van den Driessche, Thore Graepel, and Demis Hassabis. Mastering the game of Go without human knowledge. *Nature*, 550:354–359, 2017.
- [54] Richard S Sutton and Andrew G Barto. *Reinforcement learning: An introduction*. MIT press, 2018.

- [55] Paolo Toth and Daniele Vigo. *Vehicle routing: problems, methods, and applications*. SIAM, 2014.
- [56] TransitCenter. Who’s on board 2019: How to win back america’s transit riders. Technical report, TransitCenter, New York, 2019.
- [57] Transitland. Bloomington transit (bt) — operator details. <https://www.transit.land/operators/o-dnfg-bloomingtontransit>.
- [58] United Nations Department of Economic and Social Affairs. World urbanization prospects: The 2018 revision. Technical report, United Nations, 2019.
- [59] U.S. Census Bureau. Tiger/line shapefiles. <https://www.census.gov/cgi-bin/geo/shapefiles/index.php>, 2025. Accessed: 2025-10-06.
- [60] U.S. Census Bureau, Center for Economic Studies. Lehd origin-destination employment statistics (lodes). <https://lehd.ces.census.gov/data>, 2025. Accessed: 2025-10-06.
- [61] Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Lio, and Yoshua Bengio. Graph attention networks. *arXiv preprint arXiv:1710.10903*, 2017.
- [62] Evert Vermeir, Wouter Engelen, Johan Philips, and Pieter Vansteenwegen. An exact solution approach for the bus line planning problem with integrated passenger routing. *Journal of Advanced Transportation*, 2021(1):6684795, 2021.
- [63] Oriol Vinyals, Meire Fortunato, and Navdeep Jaitly. Pointer networks. *Advances in neural information processing systems*, 28, 2015.
- [64] Vukan R Vuchic. *Urban transit systems and technology*. John Wiley & Sons, 2007.
- [65] Vukan R Vuchic. *Urban transit: operations, planning, and economics*. John Wiley & Sons, 2017.
- [66] Di Weng, Ran Chen, Jianhui Zhang, Jie Bao, Yu Zheng, and Yingcai Wu. Pareto-optimal transit route planning with multi-objective monte-carlo tree search. *IEEE Transactions on Intelligent Transportation Systems*, 22(2):1185–1195, 2021.
- [67] World Bank. Promoting livable cities by investing in urban mobility. Results brief, World Bank Group, 2024.
- [68] Keyulu Xu, Chengtao Li, Yonglong Tian, Tomohiro Sonobe, Ken-ichi Kawarabayashi, and Stefanie Jegelka. Representation learning on graphs with jumping knowledge networks. In *International Conference on Machine Learning (ICML)*, 2018.
- [69] Tianmeng Yang, Yujing Wang, Zhihan Yue, Yaming Yang, Yunhai Tong, and Jing Bai. Graph pointer neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 36, pages 8832–8839, 2022.
- [70] Sunhyung Yoo, Jinwoo Brian Lee, and Hoon Han. A reinforcement learning approach for bus network design and frequency setting optimisation. *Public Transport*, 15(2):503–534, 2023.
- [71] Fang Zhao and Xiaogang Zeng. Simulated annealing–genetic algorithm for transit network optimization. *Journal of Computing in Civil Engineering*, 20(1):57–68, 2006.

## A Size of the search space

The scope of the design problem is to choose  $K = 16$  routes, each visiting 14 distinct nodes, on the Bloomington network graph  $G = (V, E)$ . Computing the exact size of the search space is computationally intensive, so we estimate it by approximating the number of simple paths (no repeated nodes) consisting of  $L = 13$  edges. The graph  $G$  has  $|V| = 143$  and  $|E| = 243$ , which gives an average degree

$$d = \frac{2|E|}{|V|} = \frac{486}{143} \approx 3.40.$$

**Random initialization.** For simple paths in sparse graphs, we approximate the number of candidate routes by accounting for the constraint that nodes cannot repeat. If each route may start from any of  $|V|$  nodes, the first step has approximately  $d$  choices, and each subsequent step has approximately  $d - 1$  choices (i.e., excluding the node just visited):

$$P_{\text{random}} \approx |V| \cdot d \cdot (d - 1)^{L-1}, \quad S_{\text{random}} = P_{\text{random}}^K.$$

Numerically,

$$(d - 1)^{12} = (2.40)^{12} \approx 3.65 \times 10^4, \quad P_{\text{random}} \approx 1.78 \times 10^7, \quad S_{\text{random}} \approx 9.5 \times 10^{115} \approx 10^{116}.$$

**Transit-center initialization.** In the experiments reported in this paper, all routes start from the transit-center hub. This removes the  $|V|$  factor from each route count:

$$P_{\text{hub}} \approx d \cdot (d - 1)^{L-1}, \quad S_{\text{hub}} = P_{\text{hub}}^K.$$

Numerically,

$$P_{\text{hub}} \approx 1.24 \times 10^5, \quad S_{\text{hub}} \approx 3.1 \times 10^{81} \approx 10^{82}.$$

Both counts are astronomically large; the transit-center-constrained space is the relevant one for our reported experiments, while the random initialization count describes the larger unconstrained initialization setting. The magnitude of these spaces makes exhaustive or near-exhaustive search infeasible, motivating alternative approaches that exploit the problem's structure. Note that this approximation does not adjust for overlapping edges between routes.

## B Frequency of Service assignment

Frequency of Service (FOS) setting is integral to the Transit Route Network Design Problem because it affects both passenger-facing performance, through waiting and in-vehicle movement times, and operator-side cost, through fleet requirements. In a sequential route-construction MDP, however, frequency decisions made for partial routes can provide unstable training targets. A two-node partial route may appear to benefit from high frequency because it serves a small local demand pair, while the same frequency may be inefficient once the route is extended and interacts with the rest of the network. We therefore assign frequencies after route construction using a deterministic max-load projection, following standard capacity-based frequency-setting principles [11, 20].

For each route  $k$ , let  $Q_{k,e}^{\text{norm}}(\Pi)$  denote the overlap-normalized segment passenger load rate, in passengers per hour, assigned to segment  $e$  of route  $k$ , and define

$$Q_{k,\max}^{\text{norm}}(\Pi) = \max_{e \in r_k} Q_{k,e}^{\text{norm}}(\Pi).$$

For this pre-simulation projection, we build an undirected, unweighted route graph from the completed route segments. Each served OD pair contributes  $\alpha D_{ij}$  trips per hour to a deterministic minimum-hop path on this graph; direct trips are paths contained on one route, while transfer trips switch routes at shared stops. Segment passenger load rates are accumulated along the selected path and then divided by the number of overlapping routes serving that segment, preventing overestimation of frequency requirements when multiple routes share the same road segment. UXsim then simulates passenger assignment, waiting, boarding, transfers, and traffic dynamics with the resulting frequencies fixed; the projection is not recomputed from realized simulated loads.

The max-load frequency projection is

$$F_k^{\text{ml}}(\Pi) = \max \left\{ 1, \left\lceil \frac{Q_{k,\max}^{\text{norm}}(\Pi)}{\delta_{\max} C_k} \right\rceil \right\}, \quad (13)$$

where  $C_k$  is bus capacity and  $\delta_{\max}$  is the maximum desired load factor. The paper's deterministic frequency vector is therefore  $\mathcal{F}(\Pi) = F^{\text{ml}}(\Pi)$ . Although the agent does not choose frequencies directly, it influences them through route construction: routes serving high-demand corridors receive higher frequencies, while overlapping routes split normalized segment load and may therefore receive lower frequencies.

**Minimality of the max-load projection.** For a fixed route set  $\Pi$ , define the capacity-feasible frequency set under fixed normalized segment loads as

$$\mathcal{F}_{\text{cap}}(\Pi) = \{F \in \mathbb{N}_{>0}^K : Q_{k,e}^{\text{norm}}(\Pi) \leq \delta_{\max} C_k F_k \quad \forall k, e \in r_k\}.$$

**Proposition B.1** (Minimal fixed-load frequency projection). *Fix a route set  $\Pi$  and normalized segment loads  $Q_{k,e}^{\text{norm}}(\Pi)$  that do not change with the frequency vector  $F$ . Assume  $C_k > 0$  for every route  $k$  and  $\delta_{\max} > 0$ . Then the max-load rule  $F^{\text{ml}}(\Pi)$  is the componentwise minimal element of  $\mathcal{F}_{\text{cap}}(\Pi)$ . Consequently, among all frequencies in  $\mathcal{F}_{\text{cap}}(\Pi)$ , it minimizes any operator-frequency cost  $C_{\text{op}}(F)$  that is componentwise nondecreasing in  $F$ .*

*Proof.* For route  $k$ , feasibility requires

$$F_k \geq \frac{Q_{k,e}^{\text{norm}}(\Pi)}{\delta_{\max} C_k} \quad \forall e \in r_k.$$

Therefore,

$$F_k \geq \frac{Q_{k,\max}^{\text{norm}}(\Pi)}{\delta_{\max} C_k}.$$

Since  $F_k$  must be a positive integer, the smallest feasible value is

$$\max \left\{ 1, \left\lceil \frac{Q_{k,\max}^{\text{norm}}(\Pi)}{\delta_{\max} C_k} \right\rceil \right\} = F_k^{\text{ml}}(\Pi).$$

The constraints separate by route, so the same argument holds for every  $k \in \{1, \dots, K\}$ . Hence any feasible  $F \in \mathcal{F}_{\text{cap}}(\Pi)$  satisfies  $F \geq F^{\text{ml}}(\Pi)$  componentwise. Any operator-frequency cost that is componentwise nondecreasing in  $F$  is therefore minimized by  $F^{\text{ml}}(\Pi)$ .  $\square$

For fixed route geometry and dispatch assumptions, the fleet requirement used in Eq. 7 is componentwise nondecreasing in route frequency, so this minimality property applies to the operator-frequency component of the simulator reward.

Proposition B.1 gives the max-load rule a precise role: for the segment loads induced by a completed route set, it is the smallest frequency vector that satisfies the fixed-load capacity condition. The full simulator reward can still depend on frequency through transfer waiting time, passenger assignment, and induced load patterns. We therefore define the value of joint frequency choices relative to this projection.

Because the route count, route length, and admissible operating choices are bounded in our experimental setting, the feasible route and frequency sets are finite. In unbounded variants, the maxima below can be replaced by suprema without changing the interpretation of the projection gap.

**Frequency-projection gap.** Let  $\mathcal{P}$  be the feasible set of route sets. For each  $\Pi \in \mathcal{P}$ , let  $\Omega_F(\Pi)$  denote a finite admissible set of positive integer frequency vectors for  $\Pi$  that contains  $F^{\text{ml}}(\Pi)$ . Let

$$J(\Pi, F) = \mathbb{E}_{\xi} [\mathcal{R}(\Pi, \mathbf{1}, F; \xi)]$$

denote the expected simulator reward for route set  $\Pi$ , fixed stop spacing  $\mathbf{1}$ , and frequency vector  $F$ . Define the joint route-frequency optimum

$$J_{\text{joint}}^* = \max_{\Pi \in \mathcal{P}} \max_{F \in \Omega_F(\Pi)} J(\Pi, F),$$

and the projected route-only optimum optimized in this work:

$$J_{\text{proj}}^* = \max_{\Pi \in \mathcal{P}} J(\Pi, F^{\text{ml}}(\Pi)).$$

The frequency-projection gap is

$$\Delta_F = \max_{\Pi \in \mathcal{P}} \left[ \max_{F \in \Omega_F(\Pi)} J(\Pi, F) - J(\Pi, F^{\text{ml}}(\Pi)) \right].$$

**Proposition B.2** (Frequency-projection bound). *The value difference between joint route-frequency optimization and the projected route-only problem satisfies*

$$0 \leq J_{\text{joint}}^* - J_{\text{proj}}^* \leq \Delta_F.$$

*Proof.* Because  $F^{\text{ml}}(\Pi) \in \Omega_F(\Pi)$ , the joint optimizer can always choose the projected frequency vector. Therefore,

$$J_{\text{joint}}^* = \max_{\Pi \in \mathcal{P}} \max_{F \in \Omega_F(\Pi)} J(\Pi, F) \geq \max_{\Pi \in \mathcal{P}} J(\Pi, F^{\text{ml}}(\Pi)) = J_{\text{proj}}^*.$$

This proves the lower bound.

For the upper bound, let  $\Pi^*$  be a route set attaining  $J_{\text{joint}}^*$ . By definition of  $\Delta_F$ ,

$$\max_{F \in \Omega_F(\Pi^*)} J(\Pi^*, F) \leq J(\Pi^*, F^{\text{ml}}(\Pi^*)) + \Delta_F.$$

Also,

$$J(\Pi^*, F^{\text{ml}}(\Pi^*)) \leq \max_{\Pi \in \mathcal{P}} J(\Pi, F^{\text{ml}}(\Pi)) = J_{\text{proj}}^*.$$

Combining the two inequalities gives

$$J_{\text{joint}}^* \leq J_{\text{proj}}^* + \Delta_F,$$

which proves the result.  $\square$

The gap  $\Delta_F$  identifies the value available from frequency choices after route geometry has been fixed. For a unit vector  $e_k$ , deliberately increasing service on route  $k$  improves a completed route set  $\Pi$  when there exists an integer increment  $q > 0$  such that

$$F^{\text{ml}}(\Pi) + qe_k \in \Omega_F(\Pi) \quad \text{and} \quad J(\Pi, F^{\text{ml}}(\Pi) + qe_k) > J(\Pi, F^{\text{ml}}(\Pi)).$$

This includes cases where a connector route has low direct demand but higher frequency improves transfer paths or changes passenger assignment. In the present study, such effects are summarized by  $\Delta_F$ ; extending AlphaTransit with a second-stage frequency policy would directly target this gap.

## C State Representation

The state ( $s_t$ ) encodes the static network and evolving design context on  $G$  with  $s_t = (X_t, \mathcal{I}, Z)$ .

**Node features** ( $X_t \in \mathbb{R}^{n \times 16}$ ). Let  $V_{\text{cur}}$  be the nodes already placed on the route under construction at time  $t$ ,  $V_{\text{cmp}}$  the union of nodes across completed routes, and  $V_{\text{core}} = V_{\text{cur}} \cup V_{\text{cmp}}$ . Let  $\mathcal{C}_t$  be the set of admissible candidates at time  $t$ , i.e., the one-hop neighbors of the current frontier that are not already in the route  $V_{\text{cur}}$ . The node features for node  $i$  can then be classified into six groups:

- Geometry and connectivity: normalized  $(x_i, y_i)$  coordinates and node degree.
- OD marginals:

$$d_{\text{out}}(i) = \sum_{j=1}^n D_{ij}, \quad d_{\text{in}}(i) = \sum_{j=1}^n D_{ji}. \quad (14)$$

- Candidate demand between  $i$  and the current route (nonzero only if  $i \in \mathcal{C}_t$ ):

$$a_{i \rightarrow \text{cur}}^{\text{cand}} = \mathbf{1}\{i \in \mathcal{C}_t\} \sum_{j \in V_{\text{cur}}} D_{ij}, \quad (15)$$

$$a_{i \leftarrow \text{cur}}^{\text{cand}} = \mathbf{1}\{i \in \mathcal{C}_t\} \sum_{j \in V_{\text{cur}}} D_{ji}. \quad (16)$$

- Designed network demand at  $i$  with respect to nodes already in any designed route (nonzero only if  $i \in V_{\text{core}}$ ):

$$a_{i \rightarrow \text{core}}^{\text{core}} = \mathbf{1}\{i \in V_{\text{core}}\} \sum_{j \in V_{\text{core}}} D_{ij}, \quad (17)$$

$$a_{i \leftarrow \text{core}}^{\text{core}} = \mathbf{1}\{i \in V_{\text{core}}\} \sum_{j \in V_{\text{core}}} D_{ji}. \quad (18)$$

- Route conditioned demand for all nodes:

$$a_{i \rightarrow \text{cur}}^{\text{all}} = \sum_{j \in V_{\text{cur}}} D_{ij}, \quad a_{i \leftarrow \text{cur}}^{\text{all}} = \sum_{j \in V_{\text{cur}}} D_{ji}, \quad (19)$$

$$a_{i \rightarrow \text{cmp}}^{\text{all}} = \sum_{j \in V_{\text{cmp}}} D_{ij}, \quad a_{i \leftarrow \text{cmp}}^{\text{all}} = \sum_{j \in V_{\text{cmp}}} D_{ji}. \quad (20)$$

- Flags: one indicator each for current route membership  $\mathbf{1}\{i \in V_{\text{cur}}\}$ , the fraction of completed routes that contain  $i$  in  $[0, 1]$ , and valid next node  $\mathbf{1}\{i \in \mathcal{C}_t\}$ .

This representation exposes the policy to demand patterns at multiple scales, from immediate next candidate nodes to the global network, enabling it to reason about ridership potential, transfers, and coverage, without requiring direct access to the OD matrix. Equations (14) summarize the baseline incoming and outgoing demands at node  $i$ , independent of the current design. Equations (15) and (16) quantify the immediate marginal gain of selecting a candidate by measuring flows between a candidate and the current route. Equations (17) and (18) summarize how well a node is integrated with the already designed network. The route conditioned terms (19) and (20) provide additional context, i.e., a global view of demand relative to the current and completed routes.

**Edge connectivity ( $\mathcal{I}$ ).** The directed edge list encodes graph topology for message passing. Since streets are bidirectional, we include both  $(u, v)$  and  $(v, u)$ , i.e.,  $\mathcal{I} = \{(u, v) \in V \times V \mid \{u, v\} \in E\}$ .

**Edge features ( $Z \in \mathbb{R}^{|\mathcal{I}| \times 2}$ ).** For each directed edge,  $Z$  stores length and free flow speed.

All continuous features are min-max scaled to  $[0, 1]$  using network level bounds. Demand aggregates are computed from  $D$  and scaled by the global reference  $\max(\max_i \sum_j D_{ij}, \max_j \sum_i D_{ji})$ . Terms that depend on  $\mathcal{C}_t$  or  $V_{\text{core}}$  are set to zero outside those sets. Binary indicators take values in  $\{0, 1\}$ .

## D Policy Network

This section provides additional architecture details for the policy-value network. The actor is permutation equivariant because it applies the same scoring function to every node, while the critic is permutation invariant because it pools node embeddings before predicting a graph-level value. For fixed hidden widths, attention heads, and block count  $B$ , the number of trainable parameters is independent of the number of nodes  $n$ .

**Shared GATv2 backbone.** At state  $s = (X, \mathcal{I}, Z)$ , the node features  $X \in \mathbb{R}^{n \times d_x}$  contain per-node spatial attributes, OD-demand aggregates, route-membership indicators, and valid-next-node indicators. The directed edge list  $\mathcal{I}$  gives graph connectivity, and  $Z \in \mathbb{R}^{|\mathcal{I}| \times 2}$  contains edge attributes, namely link length and free-flow speed.

Node features are linearly projected to 64 channels. A stack of  $B$  GATv2 [7] blocks produces intermediate node representations  $\mathbf{h}^{(b)} \in \mathbb{R}^{n \times d_b}$  for  $b \in \{1, \dots, B\}$ . Compared with GAT [61], GATv2 uses dynamic attention that depends jointly on both endpoints. Each block conditions message passing on edge attributes, so link length and free-flow speed can affect node updates.

Each block applies pre-layer normalization, a GATv2 attention layer, a nonlinear activation, feature dropout, and a residual path. The residual path is the identity when input and output widths match, and a learned linear projection otherwise. Attention dropout is applied inside the GATv2 layer. In the trained configuration, dropout modules are present but disabled.

Jumping Knowledge aggregation [68] preserves information from all message-passing depths. For each node  $v$ , the outputs of all  $B$  blocks are concatenated and projected to the final node embedding:

$$\mathbf{z}_v = W_{\text{JK}} \left[ \mathbf{h}_v^{(1)} \parallel \mathbf{h}_v^{(2)} \parallel \dots \parallel \mathbf{h}_v^{(B)} \right] + \mathbf{b}_{\text{JK}}, \quad H = [\mathbf{z}_1; \dots; \mathbf{z}_n] \in \mathbb{R}^{n \times d},$$

where  $W_{\text{JK}} \in \mathbb{R}^{d \times \sum_{b=1}^B d_b}$  is the Jumping Knowledge projection. For the default  $B = 4$  configuration, the block widths are  $[128, 128, 64, 64]$ , so JK aggregation concatenates  $128+128+64+64 = 384$  channels and projects them to  $d = 64$ . This adds  $384 \cdot 64 + 64 = 24,640$  parameters.

The default attention-head counts are  $[8, 8, 4, 4]$ . Head outputs are averaged rather than concatenated, so each block output width remains  $d_b$ . Block 1 maps 64 channels to 128 and uses a projected residual

Simulation	Policy Network		AlphaTransit		End-to-End RL			
					$\alpha=0.3$	$\alpha=1.0$		
Horizon	10,000	Node features	16	Env steps $S_{\max}$	$\approx 10^6$	Env steps	$10^6$	
Time step	1 s	Edge features	2	MCTS sims $N_{\text{iter}}$	500	Discount $\gamma$	0.999	
Bus capacity	40	GATv2 blocks	4	PUCT $c$	1.0, 1.5	GAE $\lambda$	0.95	
Stop duration	60 s	Activation	tanh	Dirichlet $\alpha_{\text{dir}}$	0.3	Value coef	0.5	
Routes $K$	16	Channels	128, 128, 64, 64	Dirichlet $\varepsilon$	0.25	Learning rate $\eta$	$5 \times 10^{-5}$	$10^{-5}$
Max length $L_{\max}$	14	Attn heads	8, 8, 4, 4	Buffer size	50k	Clip $\epsilon$	0.2	0.1
Modal split $\alpha$	0.3, 1.0	Actor MLP	256, 128, 64	Workers $W$	8 / 16	Epochs $K_e$	8	4
Access radius	0.5 km	Critic MLP	256, 128, 64	Batch size	256	Batch size	256	128
Stop spacing	1	Attn dropout	0, 0, 0, 0	Learning rate $\eta$	$10^{-4}$	Entropy coef	0.01	0.02
Platoon $\Delta n$	5	Hidden dim	64	Train steps/iter	200	LR anneal	No	Yes

Table 2: Hyperparameters for simulation, search, and training. Simulation settings are shared; policy settings apply to AlphaTransit and End-to-End RL. Pure MCTS lacks policy training and shares search budget and PUCT rule. AlphaTransit and Pure MCTS use  $N_{\text{iter}} = 500$  simulations per decision in final comparisons. Workers are 8 or 16. Sweeps selected training hyperparameters; grid sweeps varied search depth, data diversity, and model size.  $\tau$  follows  $1.0 \rightarrow 0.7 \rightarrow 0.5$ .

path; Block 2 keeps width 128 and uses an identity residual; Block 3 maps 128 channels to 64 and uses a projected residual path; Block 4 keeps width 64 and uses an identity residual.

**Actor head.** The actor uses a pointer-style scoring mechanism [63, 69]. A shared MLP scores every node embedding  $\mathbf{z}_v$  and produces one logit per node. The MLP is applied to all nodes, so the actor is permutation equivariant. A feasibility mask removes inadmissible actions, and the remaining logits are normalized to obtain  $P_\theta(\cdot | s)$ . The actor selects one node per construction step, matching the sequential route-extension process. The default actor MLP has hidden widths 256, 128, and 64.

**Critic head.** The critic applies global mean pooling and global max pooling to the node embedding matrix  $H$ , concatenates the two pooled vectors, and passes the result through an MLP to predict  $V_\theta(s)$ . Pooling makes the critic permutation invariant and produces one scalar value per graph. The default critic MLP has hidden widths 256, 128, and 64.

**Complexity and implementation.** The backbone is computed once per state to obtain  $H$ . The actor scores all nodes with a shared MLP, and the critic applies parameter-free global pooling before predicting  $V_\theta(s)$ . For fixed hidden widths, attention heads, and block count  $B$ , the trainable parameter count does not depend on  $n$ . The forward-pass cost grows with the graph as  $O(B(n + |\mathcal{I}|))$ , which is linear in  $n$  for sparse road graphs where  $|\mathcal{I}| = O(n)$ .

The input contains 16 node features and 2 edge features. The default configuration uses  $B = 4$  GATv2 blocks, final node embedding dimension  $d = 64$ , block widths [128, 128, 64, 64], and attention-head counts [8, 8, 4, 4]. More generally, for  $B$  blocks we use the channel schedule  $[128]^{[B/2]} + [64]^{[B/2]}$  and the head schedule  $[8]^{[B/2]} + [4]^{[B/2]}$ . All linear layers use orthogonal initialization, and LayerNorm parameters are initialized to unit scale and zero bias.

## E Training Procedures and Hyperparameters

Experiments were conducted on an AMD Ryzen Threadripper 7960X CPU with 377 GB RAM and an NVIDIA RTX PRO 6000 GPU. For the 1M-step runs, End-to-End RL with PPO takes roughly 3–5 hours depending on  $\alpha$ , AlphaTransit with  $N_{\text{iter}} = 500$  takes roughly 24–27 hours, and Pure MCTS takes roughly 120–168 hours. Table 2 summarizes the simulation, search, and training settings.

### E.1 Genetic Algorithm

We implement a genetic algorithm for transit network design where each individual represents a complete transit network  $\Pi = (r_1, \dots, r_K)$  of  $K$  routes, each satisfying  $L_{\min} \leq |r_k| \leq L_{\max}$  with  $L_{\min} = 2$  and  $L_{\max} = 14$ .

- Initialization: To prevent premature convergence to local optima, the initial population combines three seeding strategies: (i) demand-guided construction that greedily extends routes toward

high OD-interaction nodes, (ii) random feasible routes via constrained random walks, and (iii) a warm-start individual sampled from the existing real-world routes.

- Operators: Selection uses tournament selection with size 3. Crossover performs route exchange where each route index inherits from one parent with equal probability, preserving complete routes to maintain feasibility. Mutation applies path regeneration by cutting a route at a random interior node and regrowing via random walk, preserving promising prefixes while exploring new suffixes. A repair step ensures all routes contain at least  $L_{\min} = 2$  nodes after each operation, preventing degenerate one-node routes.
- Fitness: Each individual is evaluated by assigning frequencies via the max-load rule (Appendix B) and running simulation. Fitness uses the same terminal-only reward  $\mathcal{R}$  (Eq. 7) as AlphaTransit.

We use population size 50 with crossover rate 80%, mutation rate 40%, and elitism count 5, yielding approximately 5,000 simulator evaluations (approximately 1M environment steps) and a total wall-clock runtime of 8 hours. We run 100 generations per demand setting, with the best design found at generation 94 for  $\alpha = 0.3$  and generation 81 for  $\alpha = 1.0$ .

## E.2 Bee Colony Optimization

We adapt the implementation from Holliday et al. [24], which itself extends the original Bee Colony algorithm of Nikolić and Teodorović [41]. Two mutation operators act on the route-set population: first, replacement mutation substitutes a dropped route with a shortest-path route between randomly selected terminals; second, endpoint mutation extends or shortens an existing route by adding or removing a node at either endpoint. We use population size  $B = 15$  with  $E = 10$  mutations per iteration over 400 iterations. To match our setup, the implementation is modified so that (i) initialization places every initial route at the transit center, (ii) replacement mutations are forced to originate there, and (iii) endpoint mutations are reverted if they would remove the transit center from a route’s start. Because Bee Colony Optimization optimizes the analytical cost function  $C = C_p + C_o$ , which has no notion of modal split, a single route set is generated and then evaluated through the UXsim simulator at both  $\alpha = 0.3$  and  $\alpha = 1.0$ . The demand matrix is symmetrized as  $D'_{ij} = (D_{ij} + D_{ji})/2$  to satisfy the symmetric demand assumption of the underlying network design problem formulation.

## E.3 Neural Evolutionary Algorithm

The neural evolutionary algorithm augments Bee Colony Optimization by replacing half of the replacement mutations with route construction proposed by a graph neural network policy [23, 24]. The construction policy is trained from scratch on 32,768 synthetic 20-node cities following the PPO protocol from Holliday et al. [24], with the transit-center start constraint enforced at the model level: a mask tensor identifies the designated start node in each city, and the policy assigns  $-\infty$  logits to all other start positions so that only routes originating at that node can be produced. We retrain rather than reuse the published pre-trained weights because those weights were learned without this constraint, causing every neural mutation in our setup to be rejected and the algorithm to degrade to plain Bee Colony Optimization. The algorithm is initialized from the best of 100 samples drawn from the trained construction policy, then run for 400 iterations using the same population as Bee Colony Optimization. As with Bee Colony Optimization, a single route set is generated under the analytical objective and then evaluated through UXsim at both demand levels.

## E.4 Pure Monte Carlo Tree Search

Monte Carlo Tree Search with uniform action priors  $P(s, a) = 1/|\mathcal{C}_t|$  over admissible actions and a random rollout till terminal state followed by full route simulation for leaf evaluation. The PUCT selection rule from the main-text Eq. 10,  $N_{\text{iter}} = 500$ , and matched  $c$  values per  $\alpha$  are the same as in AlphaTransit, and a single tree is carried across route boundaries via re-rooting. Each rollout costs  $\sim 8$  s on Bloomington compared to  $\sim 1$  ms for the GNN value head; rollouts are parallelized across 8 workers, each running an independent simulator instance.

## E.5 End-to-End Reinforcement Learning

A key challenge in applying standard policy gradient methods to the Transit Route Network Design Problem (TRNDP) is the long episode horizon: with  $K = 16$  routes of up to  $L_{\max} = 14$  nodes each,

---

**Algorithm 1** End-to-End Reinforcement Learning

---

```
1: Input: Graph  $G = (V, E)$ , OD matrix  $D$ , edge index  $\mathcal{I}$ , edge features  $Z$ , routes  $K$ , max route length  $L_{\max}$ ,  
   episodes per update  $M$   
2: Output: Policy  $\pi_\theta$  yielding routes  $\Pi = (r_1, \dots, r_K)$   
3: Initialize policy-value network  $(\pi_\theta, V_\theta)$  and PPO buffer  $B \leftarrow \emptyset$   
4: while environment steps  $< S_{\max}$  do  
5:   // Parallel full-episode collection  
6:   for workers  $w = 1, \dots, N$  in parallel until  $M$  episodes are collected do  
7:     Reset environment;  $\Pi \leftarrow \emptyset$ ; initialize current route  $r_1$  from the transit center  
8:     while episode not terminated do  
9:        $s_t \leftarrow \text{FORMSTATE}(\text{current route, completed routes, } \mathcal{I}, Z)$   
10:       $\mathcal{C}_t \leftarrow$  valid one-hop frontier neighbors not already in current route  
11:       $m_t \leftarrow \text{MASK}(\mathcal{C}_t)$   
12:      if  $\mathcal{C}_t = \emptyset$  then  
13:         $a_t \leftarrow \text{NOVALIDACTION}$ ;  $\log \pi_\theta(a_t | s_t, m_t) \leftarrow 0$   
14:      else  
15:         $a_t \sim \pi_\theta(\cdot | s_t, m_t)$  and record  $\log \pi_\theta(a_t | s_t, m_t)$   
16:      end if  
17:      Execute  $a_t$  in the route-construction environment  
18:      if  $a_t$  ends a non-final route then  
19:        Add completed route to  $\Pi$ ; initialize next route from the transit center  
20:      end if  
21:      if  $a_t$  ends the final route then  
22:        Add final route to  $\Pi$ ; run UXsim once on  $\Pi$ ; set  $r_t \leftarrow \mathcal{R}$  (Eq. 7)  
23:      else  
24:        Set  $r_t \leftarrow \mathcal{R}_{\text{partial}}$  (Eq. 21), or 0 in terminal-only mode  
25:      end if  
26:      Store  $(s_t, a_t, r_t, \log \pi_\theta(a_t | s_t, m_t), V_\theta(s_t), m_t, \text{done})$  locally  
27:    end while  
28:    Compute GAE advantages  $\hat{A}_t$  and returns over the full episode; append trajectory to  $B$   
29:  end for  
30:  // PPO optimization  
31:  for epoch = 1 to  $K_{\text{epochs}}$  do  
32:    for minibatch  $b \sim B$  do  
33:       $\theta \leftarrow \theta - \eta \nabla_\theta \mathcal{L}_{\text{PPO}}(\theta)$  using Eq. 22  
34:    end for  
35:  end for  
36:  Clear  $B$  and broadcast updated weights to workers  
37: end while  
38: return  $\pi_\theta$ 
```

---

episodes can span over 200 decisions before any terminal reward is available. This creates a severe credit assignment problem [54], as the policy must learn which early decisions contributed to the final outcome. To address this, the end-to-end approach augments the reward from Eq. (7) with lightweight shaping rewards during route construction. These provide more frequent feedback to guide learning, while the simulation-based reward remains the primary signal. All reported End-to-End RL results use the same delta-based shaping rule without early-stop penalty at both demand levels. The full training procedure is summarized in Algorithm 1, where  $V_{\text{cur}}$  denotes nodes in the current route and  $V_{\text{cmp}}$  denotes nodes across all completed routes.

Helpers in Algorithm 1 refer to route-construction environment operations. FORMSTATE builds the graph observation from the current and completed routes, MASK converts  $\mathcal{C}_t$  into the binary feasibility mask used by the policy, and NOVALIDACTION denotes the forced route-finalization transition when no valid extension remains; it is not a learned stop action.

During the construction of a route  $r_k$ , the agent receives feedback based on the evolving network topology. We use the incremental change in demand coverage  $\Delta\Psi_t = \max(0, \Psi_t - \Psi_{t-1})$  to reward actions that expand the reachable O-D pairs. This delta-based formulation directly measures the value of each action rather than accumulated state, providing clearer credit assignment. The reported End-to-End RL baseline penalizes route overlap but does not include an early-stop term:

$$\mathcal{R}_{\text{partial}} = b_7 \cdot \Delta\Psi_t - b_8 \cdot \omega, \quad (21)$$

where  $\omega$  is the current route overlap ratio. We use  $b_7 = 20$  and  $b_8 = 8$ . The coverage term is marginal, while the overlap term is a step-wise exposure penalty: during PPO rollout, repeatedly overlapping partial networks accumulate overlap cost across construction steps. This shaping term is used only for the End-to-End RL baseline. AlphaTransit trains from terminal full-network evaluations, so its overlap penalty is applied through Eq. (7) once per completed route set. These shaping terms remain small relative to the simulation-based reward so that the primary learning signal still comes from passenger flow outcomes. Alternative shaping variants, including early-stop penalties, are analyzed in Fig. 1. Partial rewards are also normalized online during training. Forced route termination is treated as an environment transition, and full UXsim evaluation is performed exactly once after all  $K$  routes have been finalized.

The PPO update minimizes the negative clipped surrogate with a value loss and entropy bonus,

$$\begin{aligned} \mathcal{L}_{\text{PPO}}(\theta) &= -\mathbb{E}_t \left[ \min \left( \rho_t(\theta) \hat{A}_t, \text{clip}(\rho_t(\theta), 1 - \epsilon, 1 + \epsilon) \hat{A}_t \right) \right] + c_v L_V - c_e H, \\ \rho_t(\theta) &= \frac{\pi_\theta(a_t | s_t, m_t)}{\pi_{\theta_{\text{old}}}(a_t | s_t, m_t)}. \end{aligned} \tag{22}$$

## E.6 AlphaTransit

Training alternates between parallel data generation and network optimization. During data generation,  $W$  workers independently construct complete transit networks using Monte Carlo Tree Search guided decisions. At each step, Dirichlet noise is added to the root priors to encourage exploration:

$$P(s, a) \leftarrow (1 - \epsilon) \cdot P(s, a) + \epsilon \cdot \eta_a, \quad \eta \sim \text{Dir}(\alpha_{\text{dir}}), \tag{23}$$

where  $\alpha_{\text{dir}}$  controls noise concentration and  $\epsilon$  controls the noise weight. After ordinary route-extension actions, the search tree is re-rooted at the selected child so the corresponding subtree statistics are reused. When a route has no valid extension and is finalized automatically, we root a new tree at the forced successor state. Algorithm 2 summarizes the full AlphaTransit training loop. Training runs for  $S_{\text{max}}$  environment steps across  $W$  parallel workers using learning rate  $\eta$ .

Helpers in Algorithm 2 refer to environment operations. `INITIALIZEROUTE` starts a route at the configured transit hub, `FORMSTATE` builds the graph observation from the environment state, and `VALIDACTIONS( $s_t$ )` returns the admissible one-hop extensions  $C_t$ . `APPEND` stores each search target tuple for replay, and `ROUTES` returns the completed route tuple before terminal simulation.

- Value Normalization: Terminal rewards can vary significantly in scale across episodes. To stabilize learning, we normalize rewards online:

$$\tilde{z} = \text{clip} \left( \frac{z - \mu}{\sigma + \epsilon}, -3, 3 \right),$$

where running statistics  $(\mu, \sigma)$  are updated with all raw rewards before normalizing.

- Temperature Schedule: The temperature  $\tau$  in Eq. (11) uses a schedule based on training progress:

$$\tau(\text{progress}) = \begin{cases} 1.0 & \text{if } \text{progress} < 0.3, \\ 0.7 & \text{if } 0.3 \leq \text{progress} < 0.6, \\ 0.5 & \text{otherwise.} \end{cases}$$

During evaluation, AlphaTransit uses near-greedy selection ( $\tau = 0.1$ ) over the MCTS visit-count distribution and disables Dirichlet noise. For End-to-End RL, we use low-temperature sampling ( $\tau = 0.1$ ) rather than argmax selection. This follows standard practice in neural combinatorial optimization [29, 30, 4], where sampling is preferred over greedy selection for three reasons: (i) the policy is trained to optimize expected return under its stochastic distribution, not the argmax; (ii) deterministic selection produces identical route sets across evaluation runs, precluding any measure of variance; and (iii) AlphaTransit’s near-greedy evaluation operates over search-refined visit counts, not raw logits, so aligning PPO’s evaluation temperature provides a fairer comparison.

---

**Algorithm 2** AlphaTransit

---

```
1: Input: Graph  $G$ , OD matrix  $D$ , edges  $\mathcal{I}$ , features  $Z$ , routes  $K$ , max length  $L_{\max}$ , MCTS
simulations  $N_{\text{iter}}$ , workers  $W$ .
2: Output: Policy-value network  $f_\theta$  yielding route tuples  $\Pi = (r_1, \dots, r_K)$ .
3: Initialize: Network  $f_\theta$ , replay buffer  $\mathcal{D} \leftarrow \emptyset$ , reward statistics  $\mathcal{Z}$ .
4: while  $t_{\text{env}} < S_{\max}$  do
5:   Broadcast current parameters  $\theta$  to workers; set temperature  $\tau$ .
6:   // Parallel MCTS-guided episode collection
7:   for each worker  $w \in \{1, \dots, W\}$  in parallel do
8:     Reset environment;  $r_1 \leftarrow [\text{INITIALIZEROUTE}()]$ ; initialize MCTS tree  $\mathcal{T}$ .
9:      $\mathcal{E}_w \leftarrow []$ ;  $H_w \leftarrow 0$ .
10:    while route set is incomplete do
11:       $s_t \leftarrow \text{FORMSTATE}(\mathcal{I}, Z)$ ;  $\mathcal{C}_t \leftarrow \text{VALIDACTIONS}(s_t)$ .
12:      if  $\mathcal{C}_t = \emptyset$  then
13:        Finalize current route; initialize the next route if one remains; reset  $\mathcal{T}$ .
14:         $H_w \leftarrow H_w + 1$ ; continue.
15:      end if
16:      Run  $N_{\text{iter}}$  MCTS simulations from  $s_t$  using  $f_\theta$  and PUCT.
17:      Derive  $\pi_t(a | s_t) \propto N(s_t, a)^{1/\tau}$  over  $a \in \mathcal{C}_t$ .
18:       $\mathcal{E}_w.\text{APPEND}(s_t, \pi_t, \mathcal{C}_t)$ .
19:      Sample  $a_t \sim \pi_t$ ; apply  $a_t$ ; re-root  $\mathcal{T}$  at the selected child.
20:       $H_w \leftarrow H_w + 1$ .
21:    end while
22:     $\Pi_w \leftarrow \text{ROUTES}()$ ;  $z_w \leftarrow R_T(\Pi_w; \xi)$ .
23:  end for
24:  // Aggregate episodes
25:  for  $w = 1$  to  $W$  do
26:    Update reward statistics  $\mathcal{Z}$  with raw reward  $z_w$ .
27:    Add  $(s_i, \pi_i, \mathcal{C}_i, z_w)$  to  $\mathcal{D}$  for all  $(s_i, \pi_i, \mathcal{C}_i) \in \mathcal{E}_w$ .
28:  end for
29:   $t_{\text{env}} \leftarrow t_{\text{env}} + \sum_w H_w$ .
30:  // Policy-value network optimization
31:  for step = 1 to  $N_{\text{steps}}$  do
32:    Sample minibatch  $\mathcal{B} = \{(s_j, \pi_j, \mathcal{C}_j, z_j)\}_{j=1}^m \subseteq \mathcal{D}$ .
33:    Normalize each  $z_j$  with  $\mathcal{Z}$  to obtain  $\tilde{z}_j$ .
34:    Evaluate  $f_\theta(s_j)$  and mask  $P_\theta$  over  $\mathcal{C}_j$  for each tuple.
35:    Define  $\ell_j(\theta) = -\sum_{a \in \mathcal{C}_j} \pi_j(a | s_j) \log P_\theta(a | s_j) + (V_\theta(s_j) - \tilde{z}_j)^2$ .
36:     $\theta \leftarrow \theta - \eta \nabla_\theta \frac{1}{m} \sum_{j=1}^m \ell_j(\theta)$ .
37:  end for
38: end while
39: return  $f_\theta$ .
```

---

## F Real-world Networks and Data

### F.1 Bloomington, Indiana

We introduce a novel city-scale transit network dataset for Bloomington, Indiana. Unlike existing benchmark networks from literature, our dataset uniquely captures real-world aspects in three dimensions: (i) the underlying transportation network, (ii) travel demand derived from census data, and (iii) transit routes currently operating in the city.

**Network Structure.** The network consists of 143 nodes and 243 bidirectional edges, covering an area of approximately 152.3 km<sup>2</sup>. The network topology was derived from road infrastructure with several practical assumptions made to balance modeling fidelity with simulation efficiency.

- **Planar Representation:** Three-dimensional infrastructure elements such as tunnels, overpasses, and underpasses are modeled as planar connections.

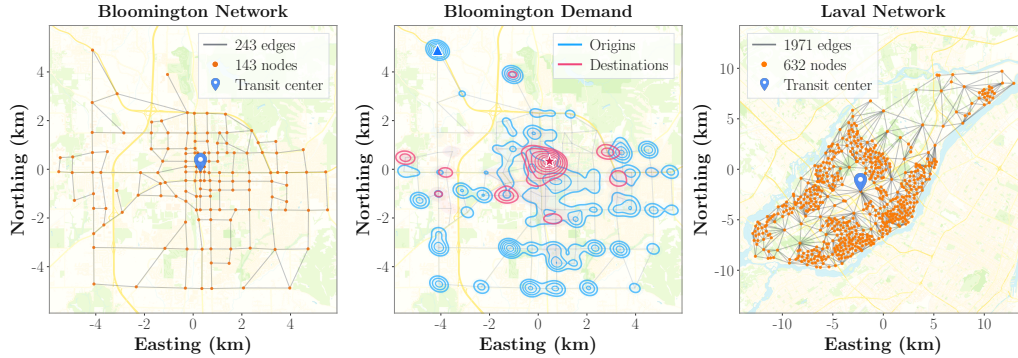


Figure 7: The Bloomington, Indiana transportation network ( $\sim 152.3 \text{ km}^2$ ). The demand map shows the spatial distribution of trip origins (blue) and destinations (red), with peak origin demand of 344 trips per hour at node 1 ( $\triangle$ ) and peak destination demand of 1,681 trips per hour at node 129 ( $\star$ ). The Laval, Quebec transportation network ( $\sim 256 \text{ km}^2$ ) is used as the out-of-distribution generalization benchmark. The Laval network, with 632 nodes and 1,971 links, is obtained from Holliday et al. [24]; the blue transit-center marker indicates node 542.

- **Edge Geometry:** All edges are represented as bidirectional. When two parallel one-way streets exist next to each other, they are consolidated into a single bidirectional edge positioned at their centerline. Further, edges follow shortest-distance connections between nodes rather than exact street curvatures; the length differences are negligible for the scale of analysis.
- **Speed Assignment:** All edges are assigned a uniform free-flow speed of 16.67 m/s (60 km/h or 37 miles/hr), reflecting typical urban traffic speed limit.
- **Highway Exclusion:** Interstate 69 highway segments within the city limits were excluded from the transit network, as city transit primarily serves local destinations and highways lack appropriate passenger access points. The bus routes currently operating in the city also avoid the highway.
- **Shared Routes:** Although existing bus routes may have slightly different outbound and inbound paths, to reduce complexity, we model them as identical paths, manually selecting the most appropriate nodes (based on access, length, and community served) to maintain essential connectivity.

**Coordinate Transformation.** The raw geospatial data uses geographic coordinates (latitude and longitude) in angular units. However, operations in our processing pipeline such as calculating census block centroids and edge distances require a flat, Cartesian plane for accurate results. Therefore, the coordinates were transformed to a Cartesian coordinate system using the Universal Transverse Mercator (UTM) Zone 16N projection, which covers Indiana including Bloomington. The UTM projection preserves local angles and shapes while providing true metric distances and areas with minimal distortion, making it ideal for our case.

**Demand Generation.** The primary demand component was obtained from commuting trips captured in the 2022 LEHD Origin-Destination Employment Statistics (LODES) from the U.S. Census Bureau [60], which provides flows between census blocks with home locations as origins and work locations as destinations. We utilize block-level census data [59] and processed a total of 2,399 census blocks within Monroe County, Indiana (FIPS code 105, GEOID 18105), which was further reduced to 1,475 blocks to confine the area of interest to the vicinity of Bloomington. For each census block, its centroid was calculated and the demand origins and destinations in that block were assigned to the node nearest to the centroid.

Because the LODES data excludes trips for non-commuting purposes such as school and shopping (typically classified as home-based other or non-home-based in transportation research), to account for this mixed composition of traffic, we scale the commuting flows by 150%, consistent with typical values ranging from 100% – 200% depending on time of day [37]. Further, to express demand on an hourly basis, we adopt a peak-hour share of 11% of daily traffic, which lies within the typical 6% – 12% range [44]. The resulting origin–destination (OD) demand matrix contains 5,737 pairs, with a maximum origin demand of 344 trips per hour and a maximum destination demand of 1,681 trips per hour. This OD matrix is treated as an exogenous peak-hour trip table. In each experiment,  $\alpha$  is fixed before evaluation and scales the table into a transit-demand scenario.




Figure 8: Learning dynamics and search cost under full demand ( $\alpha = 1.0$ ). **LEFT**: Curves average two training seeds per reward mode. Reward shaping remains important for End-to-End RL, although the higher service-weight objective is more forgiving than under mixed demand ( $\alpha = 0.3$ ). Early Stopping (ES) penalizes routes that terminate before  $L_{\max}$ , while Delta-coverage ( $\Delta$ ) rewards only newly covered demand. The  $\Delta$  variants give the strongest final performance, with  $\Delta + ES$  ending slightly highest. **MIDDLE**: MCTS-provided search targets improve sample efficiency under the same environment-step budget. At search depth  $N_{\text{iter}} = 500$ , AlphaTransit leads from earliest plotted evaluations and finishes with reward 39.93, compared with 27.92 for End-to-End RL. **RIGHT**: Learned policy and value estimates make MCTS search practical across search depths  $N_{\text{iter}} \in [100, 500]$  on the Bloomington network using one CPU worker. At  $N_{\text{iter}} = 500$ , AlphaTransit requires 6.56 seconds per decision, while Pure MCTS requires 695.48 seconds per decision.

**Existing Bus Routes.** The Bloomington Transit system [6, 57] operates 16 bus routes which map to an average of 14.2 nodes per route in our network (ranging from 8 to 24 nodes).

## F.2 Laval, Quebec

To evaluate cross-city generalization, we use the Laval, Québec network from Holliday et al. [24]. The Laval graph contains 632 nodes and 1,971 bidirectional edges over approximately 256 km<sup>2</sup>, making it roughly 4.4× larger than Bloomington in node count. We use node 542 as the transit center because it has the highest degree of 21, highest closeness centrality, and highest betweenness centrality in the network. Despite the larger node count, Laval has the same graph diameter of 17 hops and a comparable mean shortest-path length of 7.1 hops because of its higher connectivity, with average degree 6.24 compared with 3.40 for Bloomington. This structural similarity supports using the same route-design parameters as in Bloomington, namely  $K=16$  routes and maximum route length  $L_{\max}=14$ . In Laval, 99.7% of node pairs are reachable within 14 hops, so this length bound remains appropriate despite the larger graph. For simulation, Laval’s unscaled demand rate is about 548K trips per hour, roughly 60× Bloomington’s, and is uniformly scaled by 0.14× to match per-link traffic density, yielding about 77K trips per hour. The Bloomington-trained agent is evaluated directly on Laval without additional training or fine-tuning. Because the shared Laval dataset does not include real-world bus routes, the Real-World baseline is excluded from this evaluation.

## G Extended Results

We evaluate every transit design through the same simulation pipeline. At simulation end,  $N_{\text{comp}}$ ,  $N_{\text{ongoing}}$ , and  $N_{\text{waiting}}$  denote passengers who completed trips, are onboard buses, and are waiting at stops, respectively. We set  $N_{\text{boarded}} = N_{\text{served}} = N_{\text{comp}} + N_{\text{ongoing}}$  for passengers counted as served, and  $N_{\text{want}}$  denotes potential riders. We use  $\tau_p^{\text{wait}}$  and  $\tau_p^{\text{move}}$  for passenger-level accumulated waiting and in-vehicle movement times, respectively, and variables with an overbar denote averages in minutes. The seven evaluation metrics are:

- Service rate (%): Fraction of potential demand counted as served,  $\sigma = N_{\text{boarded}}/N_{\text{want}}$ . This differs from the fixed-denominator service term  $\rho = N_{\text{boarded}}/N_{\text{OD}}$  used by the training reward.
- Wait time (min): Average waiting time over served riders,  $\bar{t}_{\text{wait}}$ .
- Transfer rate (%): Fraction of completed trips requiring at least one transfer,  $N_{\text{transfer}}/N_{\text{comp}} \times 100$ .

Method	Passenger Metrics				Operator Metrics			
	Service Rate (%) $\uparrow$	Wait Time (min) $\downarrow$	Transfer Rate (%) $\downarrow$	Journey Time (min) $\downarrow$	Route Efficiency $\uparrow$	Fleet Size $\downarrow$	Bus Util. (%) $\uparrow$	
$\alpha = 0.3$	Real-World	42.77 $\pm$ 1.05	14.02 $\pm$ 0.56	86.05 $\pm$ 0.44	46.83 $\pm$ 1.09	13.15 $\pm$ 0.42	89.00	17.69 $\pm$ 0.18
	Random Walk	36.41 $\pm$ 3.68	21.79 $\pm$ 2.75	82.26 $\pm$ 3.39	49.66 $\pm$ 3.52	5.06 $\pm$ 1.27	35.40 $\pm$ 4.54	11.10 $\pm$ 2.10
	Demand Cover	40.11 $\pm$ 3.43	19.15 $\pm$ 3.52	77.45 $\pm$ 3.29	45.46 $\pm$ 4.61	6.18 $\pm$ 1.76	40.10 $\pm$ 10.95	12.42 $\pm$ 2.74
	Shortest Path	37.89 $\pm$ 4.17	23.35 $\pm$ 2.73	84.43 $\pm$ 2.56	45.06 $\pm$ 4.06	4.75 $\pm$ 0.80	21.60 $\pm$ 2.58	6.53 $\pm$ 1.07
	Genetic Alg.	50.42 $\pm$ 0.50	9.07 $\pm$ 0.50	82.61 $\pm$ 0.27	39.14 $\pm$ 0.59	17.12 $\pm$ 0.21	79.00	19.76 $\pm$ 0.21
	Bee Colony	39.83 $\pm$ 0.66	10.50 $\pm$ 0.49	87.97 $\pm$ 0.20	34.09 $\pm$ 0.44	18.58 $\pm$ 0.25	94.00	12.07 $\pm$ 0.09
	Neural Evol.	47.85 $\pm$ 0.57	5.65 $\pm$ 0.20	92.16 $\pm$ 0.11	29.94 $\pm$ 0.29	23.71 $\pm$ 0.22	101.00	13.11 $\pm$ 0.16
	Pure MCTS	53.30 $\pm$ 0.97	7.42 $\pm$ 0.54	84.97 $\pm$ 2.58	37.57 $\pm$ 1.34	19.87 $\pm$ 1.01	86.00 $\pm$ 4.15	21.93 $\pm$ 1.26
	End-to-End RL	49.72 $\pm$ 1.77	8.46 $\pm$ 1.28	84.49 $\pm$ 1.08	41.14 $\pm$ 1.77	18.09 $\pm$ 1.43	111.90 $\pm$ 6.04	19.89 $\pm$ 1.02
	AlphaTransit	54.64 $\pm$ 0.54	7.13 $\pm$ 0.16	82.66 $\pm$ 0.29	35.81 $\pm$ 0.40	21.99 $\pm$ 0.33	80.00	22.10 $\pm$ 0.18
$\alpha = 1.0$	Real-World	58.44 $\pm$ 0.95	15.95 $\pm$ 0.39	81.90 $\pm$ 0.31	52.85 $\pm$ 0.72	61.83 $\pm$ 0.92	281.00	31.89 $\pm$ 0.25
	Random Walk	62.79 $\pm$ 5.55	15.72 $\pm$ 2.86	79.87 $\pm$ 2.66	47.45 $\pm$ 3.22	36.85 $\pm$ 9.30	109.40 $\pm$ 31.86	28.61 $\pm$ 3.77
	Demand Cover	58.03 $\pm$ 7.49	16.28 $\pm$ 2.61	75.08 $\pm$ 4.23	47.49 $\pm$ 3.26	35.52 $\pm$ 7.52	124.60 $\pm$ 32.86	26.83 $\pm$ 2.83
	Shortest Path	56.50 $\pm$ 12.49	20.37 $\pm$ 4.73	69.60 $\pm$ 8.07	42.84 $\pm$ 2.88	23.47 $\pm$ 8.92	48.00 $\pm$ 12.20	18.15 $\pm$ 3.34
	Genetic Alg.	81.17 $\pm$ 0.99	8.94 $\pm$ 0.45	87.28 $\pm$ 0.23	44.32 $\pm$ 0.57	100.19 $\pm$ 1.87	254.00	43.16 $\pm$ 0.29
	Bee Colony	64.74 $\pm$ 1.49	11.96 $\pm$ 0.60	84.45 $\pm$ 0.20	48.58 $\pm$ 0.67	89.36 $\pm$ 0.98	301.00	28.78 $\pm$ 0.48
	Neural Evol.	70.51 $\pm$ 1.85	11.91 $\pm$ 0.91	87.34 $\pm$ 0.23	49.20 $\pm$ 0.51	99.64 $\pm$ 1.42	320.00	30.39 $\pm$ 0.49
	Pure MCTS	73.79 $\pm$ 4.96	8.68 $\pm$ 1.91	80.07 $\pm$ 2.66	43.92 $\pm$ 3.22	77.39 $\pm$ 5.41	200.80 $\pm$ 24.75	36.29 $\pm$ 1.81
	End-to-End RL	73.70 $\pm$ 1.68	10.04 $\pm$ 0.82	85.29 $\pm$ 0.53	50.52 $\pm$ 1.30	83.78 $\pm$ 4.40	346.50 $\pm$ 14.54	43.31 $\pm$ 1.57
	AlphaTransit	82.08 $\pm$ 0.55	8.48 $\pm$ 0.39	82.55 $\pm$ 0.19	43.10 $\pm$ 0.49	110.58 $\pm$ 0.90	267.00	45.02 $\pm$ 0.24

Table 3: Performance comparison across ten methods under mixed demand ( $\alpha = 0.3$ ) and full transit demand ( $\alpha = 1.0$ ) on the Bloomington benchmark. AlphaTransit achieves the highest service rate in both demand settings, with 54.64% under mixed demand and 82.08% under full transit demand; under full transit demand, it also achieves the best wait time, route efficiency, and bus utilization. Arrows indicate the direction of improvement ( $\uparrow$  higher is better,  $\downarrow$  lower is better), and values report mean  $\pm$  standard deviation over 10 evaluation seeds for each trained policy or generated route set. AlphaTransit uses  $N_{\text{iter}} = 500$ , while End-to-End RL is the PPO baseline without MCTS. Fleet-size deviations appear only when the route-design procedure yields different route sets across seeds, since frequencies and fleet size are deterministic once a route set is fixed.

- Journey time (min): Average elapsed passenger journey time for passengers counted as served,

$$\bar{t}_{\text{journey}} = \frac{1}{N_{\text{served}}} \sum_{p \in \mathcal{S}_{\text{served}}} (\tau_p^{\text{wait}} + \tau_p^{\text{move}}), \quad (24)$$

where  $\mathcal{S}_{\text{served}}$  contains passengers who completed trips or are onboard at the simulation end. This metric includes both waiting and in-vehicle movement.

- Route efficiency (pax/km): Passengers served per unit infrastructure,  $N_{\text{comp}}/L_{\text{total}}$ , where  $L_{\text{total}}$  is the total route length in kilometers.
- Fleet size: Total number of buses deployed across all routes,  $N_{\text{bus}}$ .
- Bus utilization (%): Average load across all buses,  $u$ , computed as occupancy divided by capacity.

The demand coverage potential  $\Psi$  is defined in Section 3. The route overlap ratio  $\omega$  is computed over unique undirected road segments. Let  $\mathcal{E}_{\Pi}$  be the set of segments used by at least one route, let  $c_e$  be the number of nonempty routes containing segment  $e$ , and let  $K_{\text{eff}}$  be the number of routes with at least one segment. Each route contributes a segment at most once, so duplicate traversals within one route are not double-counted. We define

$$\omega(\Pi) = \begin{cases} \frac{1}{|\mathcal{E}_{\Pi}|} \sum_{e \in \mathcal{E}_{\Pi}} \frac{c_e - 1}{K_{\text{eff}} - 1}, & K_{\text{eff}} > 1, \\ 0, & K_{\text{eff}} \leq 1. \end{cases} \quad (25)$$

Thus  $\omega = 0$  when no road segment is shared across routes and  $\omega = 1$  when every used segment appears in every nonempty route. This is an edge-count-based measure, not a length-weighted or node-overlap measure.

Table 3 gives the full metric-level comparison behind the main-text results. AlphaTransit achieves the strongest service-rate outcome at both modal splits, reaching 54.64% at  $\alpha = 0.3$  and 82.08% at




Figure 9: Full transit demand results on the Bloomington benchmark ( $\alpha = 1.0$ ). Points show means; error bars denote  $\pm 1$  standard deviation. In each panel, green overlay and Optimal label mark the direction of improvement: upper-left for service rate versus fleet size, lower-left for journey time versus transfer rate, and upper-right for bus utilization versus route efficiency. **LEFT**: AlphaTransit achieves highest service rate, 82.08%, with fleet size 267. **MIDDLE**: AlphaTransit obtains 43.10 minutes of journey time and an 82.55% transfer rate. **RIGHT**: AlphaTransit achieves highest bus utilization, 45.02%, and highest route efficiency, 110.58. Together, the panels show that AlphaTransit pairs the highest service rate and operator efficiency with a competitive passenger-burden trade-off.

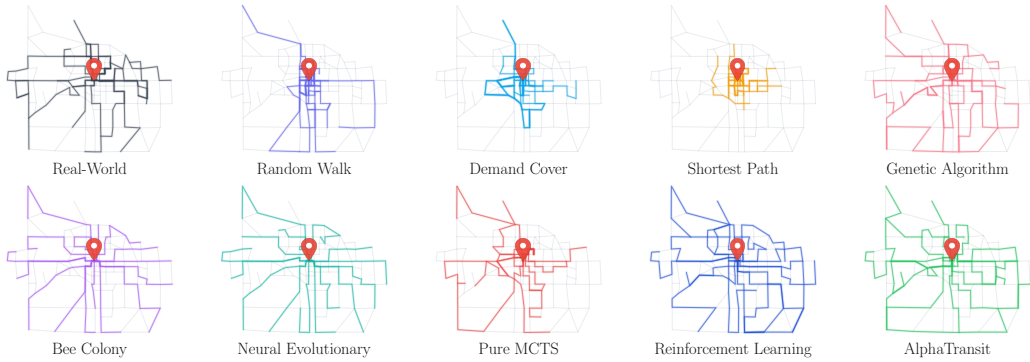


Figure 10: Selected route designs under full transit demand ( $\alpha = 1.0$ ) on the Bloomington street network. Each panel overlays selected routes on the same street basemap and marks the transit center. Pure MCTS denotes search without a learned policy/value network, while AlphaTransit uses neural network-guided MCTS. Structural-analysis values show AlphaTransit serves 128 nodes, corresponding to 89.5% node coverage, with 20.5% shared-edge overlap and route distance 129.0 km. In comparison, Real-World routes serve 114 nodes with 79.7% node coverage and 34.8% shared-edge overlap, while End-to-End RL serves 135 nodes with 94.4% node coverage and 19.3% shared-edge overlap i.e., AlphaTransit is not simply maximizing node coverage relative to End-to-End RL.

$\alpha = 1.0$ . The  $\alpha = 0.3$  setting highlights the trade-off structure: AlphaTransit has the highest service rate and bus utilization, while Neural Evolutionary has lower wait times and journey times at lower served demand and Shortest Path uses the smallest fleet. At  $\alpha = 1.0$ , AlphaTransit is stronger across most passenger and operator metrics, with the best service rate, wait time, route efficiency, and bus utilization. Relative to End-to-End RL, its service rate is 9.9% and 11.4% higher across the two regimes; relative to Pure MCTS, it is 2.5% and 11.2% higher.

Fig. 8 shows that the same qualitative trends hold under the  $\alpha = 1.0$  objective. End-to-End RL benefits from dense reward shaping, with the  $\Delta$ -coverage variants giving the strongest PPO learning curves, but AlphaTransit still achieves higher reward under the same environment-step budget. The search-time panel further shows that the learned policy/value network reduces MCTS decision cost by roughly two orders of magnitude relative to Pure MCTS at the same search depth, making search-guided route construction practical.

## **H Broader Impacts**

Urban population is projected to reach 60% of the global population by 2030 [58], placing pressure on transit infrastructure supporting 7.7 billion annual trips in the United States alone [3]. In many large cities, less than 20% of jobs can be reached within an hour using affordable transit, disproportionately affecting low-income workers [67]. AlphaTransit is intended as decision support for transit agencies that must improve service quality under fleet, budget, and planning constraints, including frequency and wait time metrics central to ridership [56]. By evaluating candidate route networks with realistic topology, census-derived demand, and simulation-based outcomes, the framework can help agencies compare coverage, wait time, completed trips, and fleet use within a single evaluation loop. Potential harms could arise if optimized plans replace community input or if aggregate demand objectives underserve riders whose needs are not well captured in the data. Automated planning should therefore complement stakeholder engagement, and future work should incorporate fairness and accessibility constraints so algorithmic transit design supports inclusive mobility goals.