

000 NATIVE ADAPTIVE SOLUTION EXPANSION FOR 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 NATIVE ADAPTIVE SOLUTION EXPANSION FOR 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 DIFFUSION-BASED COMBINATORIAL OPTIMIZATION

Anonymous authors

Paper under double-blind review

ABSTRACT

One central challenge in Neural Combinatorial Optimization (NCO) is handling hard constraints efficiently. Beyond the two classic paradigms, i.e., Local Construction (LC), which sequentially builds feasible solutions but scales poorly, and Global Prediction (GP), which produces one-shot heatmaps yet struggles with constraint conflicts, the recently proposed Adaptive Expansion (AE) shares the advantages of both by progressively growing partial solutions with instance-wise global awareness. However, existing realizations bolt AE onto external GP predictors, so their solution quality is bounded by the backbone and their inference cost scales with repeated global calls. In this paper, we fundamentally rethink adaptive expansion and make it native to a generative model, acting as its intrinsic decoding principle rather than an external wrapper. We propose NEXCO, a CO-specific masked diffusion framework that turns adaptive expansion into the model’s own iterative unmasking process. Specifically, it involves a solution-expansion training procedure with a time-agnostic GNN denoiser, which learns diffusion trajectories between fully masked solutions and ground-truth solutions. With the trained time-agnostic denoiser, we introduce a novel solution expansion scheme at the solving stage, enabling adaptive control over the intermediate solution states. It is achieved by constructing candidate sets according to confidence scores and applying feasibility projection to expand the solution while respecting constraints. In this way, “adaptive” is not an afterthought but the decoding itself: intermediate diffusion states are meaningful partial solutions and progress is instance-adaptive rather than schedule-bound. Extensive experiments on representative CO problems show that NEXCO achieves approximately 50% improvement in solution quality and up to 4 \times faster inference compared to prior state-of-the-art solvers.

1 INTRODUCTION

Combinatorial optimization (CO) is a sub-filed of mathematical optimization that involves finding the optimal solution from the discrete feasible sets. Due to their inherent NP-hardness, solving large-scale instances efficiently remains a longstanding challenge. Recent progress in Neural Combinatorial Optimization (NCO) has reduced reliance on handcrafted heuristics by learning data-driven solvers (Bengio et al., 2021; Qiu et al., 2022; Sun & Yang, 2023; Li et al., 2024; Ma et al., 2025).

Existing neural constructive solvers fall into two extreme paradigms: (1) *Local Construction (LC)* methods sequentially generates solutions in an autoregressive manner, ensuring feasibility but suffering from myopic decisions and poor scalability (Kool et al., 2019; Kwon et al., 2020; Berto et al., 2023; Drakulic et al., 2023; Pan et al., 2025); (2) *Global Prediction (GP)* methods predict full probability heatmaps in one shot, capturing global structure efficiently but producing smooth distributions that cause noisy decoding and constraint violations (Joshi et al., 2019; Fu et al., 2021; Qiu et al., 2022; Min et al., 2023; Sun & Yang, 2023; Xia et al., 2024; Li et al., 2024; Xiao et al., 2024). To bridge LC and GP, the Adaptive Expansion (AE) paradigm was introduced (Ma et al., 2025), which adaptively determines the number of variables per step. While effective, current implementations (e.g., COExpander) are merely wrappers around GP backbones such as Fast T2T, leaving two fundamental issues: (i) performance is bounded by the backbone predictor, and (ii) inference cost scales as $O(D_s \cdot C_{GP})$, far higher than vanilla GP solvers. Table 1 provides a structured summary of these trade-offs, contrasting LC, GP, and AE with our proposed method across multiple dimensions.

054
 055
 056
 057
 058
Table 1: Comparison of NCO paradigms. Local Construction (LC) and Global Prediction (GP)
 represent two extremes; COExpander realizes Adaptive Expansion (AE) as a costly wrapper around
 GP, adjusting decision granularity externally; NEXCO makes AE *native*, embedding adaptive ex-
 pansion directly into a CO-specific masked diffusion for efficient, feasible solution construction.
 NFEs: number of function evaluations; T_s : diffusion inference steps; D_s : AE expansion steps.

	LC	GP	AE	NEXCO (Ours)
Decoding granularity	One-by-one	All-at-once	Adaptive (wrapper)	Adaptive (native)
Partial feasibility	✓	✗	✓	✓
Global awareness	✗	✓	✓	✓
Complexity (NFEs)	$O(n)$	$O(T_s)$	$O(D_s * T_s)$	$O(T_s)$

064
 065 This raises a natural question: **can adaptive expansion be made native, i.e., encoded as the intrinsic**
 066 **decoding principle of a generative model, rather than an external wrapper?** Here, native AE
 067 would mean: (i) expansion progress and step size are instance-adaptive, driven by model confidence
 068 and constraints, without relying on fixed timestep schedules or external GP predictors; (ii) interme-
 069 diate states are valid, constraint-aware partial solutions, where variable commitments are enforced
 070 through feasibility projection rather than deferred to post-hoc heuristic heatmap decoding. Diffusion
 071 models appear promising: their iterative refinement resembles constructive expansion. However, ex-
 072 isting diffusion-based solvers (e.g., DIFUSCO, T2T/Fast T2T (Sun & Yang, 2023; Li et al., 2023;
 073 2024)) still operate in the GP paradigm: they generate global probability heatmaps and rely on
 074 heuristic decoding, without effectively leveraging intermediate states as partial solutions. We argue
 075 that two factors hinder this: (i) intermediate states lack semantic meaning as partial solutions, and
 076 (ii) fixed timestep schedules rigidly control denoising progress, preventing instant-adaptive progress.

077 In parallel, the broader generative modeling community has advanced Masked Diffusion Models
 078 (MDMs) as a powerful alternative to autoregressive decoding in large language models (Austin et al.,
 079 2021; Sahoo et al., 2024; Ou et al., 2025; Zheng et al., 2025; Nie et al., 2025b), which progressively
 080 unmask tokens instead of denoising noise. MDMs provide meaningful intermediate states, allow
 081 schedule-free training, and support efficient parallel decoding. *These properties closely match the*
 082 *needs of NCO, making masked diffusion a natural foundation for native adaptive expansion.*

083 Built on this insight, we propose **NEXCO** (Native Adaptive Expansion for Combinatorial
 084 Optimization), a CO-specific masked diffusion framework that embeds adaptive expansion as the
 085 intrinsic decoding principle. During training, a time-agnostic GNN denoiser learns to reconstruct
 086 ground-truth solutions from corrupted partial states, enforcing optimization consistency across dif-
 087 ferent noise levels. At inference, NEXCO adopts a Native Adaptive Expansion (NAE) strategy: the
 088 denoiser produces confidence scores, candidate sets are formed accordingly, and feasibility projec-
 089 tion ensures that selected variables satisfy problem constraints. Through iterative refinement, partial
 090 solutions progressively evolve into complete feasible ones. In contrast to COExpander (Ma et al.,
 091 2025), which implements AE as an external wrapper around global predictors, NEXCO integrates
 092 it directly into the diffusion process, achieving the efficiency of GP, the feasibility of LC, and the
 093 adaptivity of AE in a unified framework. **The contribution of this paper are:**

094 1) We revisit the Adaptive Expansion (AE) paradigm and point out that existing implementations
 095 (e.g., COExpander (Ma et al., 2025)) are merely wrappers around GP predictors, leaving their per-
 096 formance bounded by the backbone and their inference complexity scaling as $O(D_s \cdot C_{GP})$.

097 2) We propose **NEXCO**, a masked diffusion framework that natively realizes the AE paradigm by
 098 coupling a CO-specific corruption process with a time-agnostic denoiser and introducing the Native
 099 Adaptive Expansion (NAE) inference strategy. In doing so, NEXCO embeds feasibility, global
 100 awareness, and adaptivity directly into the generative process.

101 3) Extensive experiments on representative CO problems (MIS, TSP, CVRP) show that NEXCO
 102 consistently surpasses prior state-of-the-art solvers in both solution quality and inference efficiency.

103 2 PRELIMINARIES AND RELATED WORK

104 2.1 COMBINATORIAL OPTIMIZATION ON GRAPHS

105 Following standard formulations in neural CO (Sun & Yang, 2023; Li et al., 2023; 2024; Ma et al.,
 106 2025), we represent a problem instance as a graph $G(V, E)$, where V and E denote node and edge
 107 sets, respectively, and let $n = |V|$ denote the number of nodes. Decision variables are binary vectors

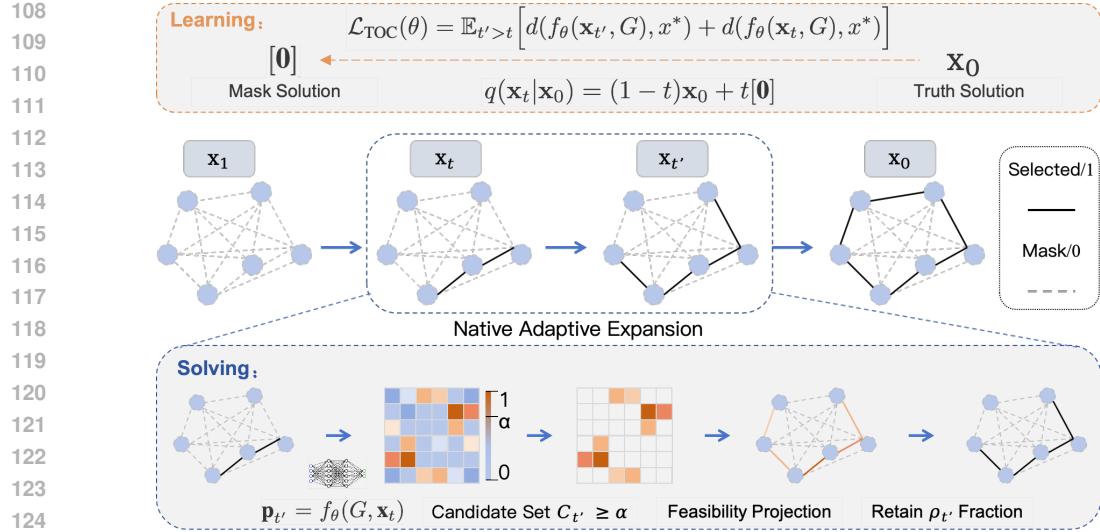


Figure 1: **Overview of NEXCO.** Given a graph instance G , a ground-truth solution $\mathbf{x}_0 \in \{0, 1\}^N$ is corrupted by masking only selected variables (1s) to 0s while keeping unselected ones (0s), producing a partial solution \mathbf{x}_t . A time-agnostic GNN denoiser f_θ predicts confidence scores for all variables without timestep conditioning. During inference, NEXCO performs *Native Adaptive Expansion* (NAE): starting from a fully masked solution $\mathbf{x}_1 = [\mathbf{0}]$, the model progressively unmasks high-confidence variables while a problem-specific projector $\Gamma(\cdot)$ enforces feasibility. This process yields valid intermediate partial solutions and converges to a complete feasible solution.

$\mathbf{x} \in \{0, 1\}^N$: for *edge-selection* problems, $N = n^2$ and $\mathbf{x}_{i \cdot n + j} = 1$ indicates whether edge (i, j) is selected; for *node-selection* problems, $N = n$ and $\mathbf{x}_i = 1$ indicates whether node i is selected. The feasible region Ω encodes hard constraints, and the objective is

$$\min_{\mathbf{x} \in \{0, 1\}^N} l(\mathbf{x}; G) \quad \text{s.t. } \mathbf{x} \in \Omega. \quad (1)$$

We study three canonical NP-hard tasks: **TSP**: find a minimum-weight Hamiltonian cycle in a complete graph; **MIS**: find a maximum-cardinality independent set; **CVRP**: minimize routing cost subject to degree and vehicle-capacity constraints.

2.2 DIFFUSION SOLVERS FOR CO

Diffusion models define a forward corruption $q(\mathbf{x}_t | \mathbf{x}_{t-1})$ and a reverse denoising $\mathbf{p}_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t)$. For binary CO, a natural adaptation is *uniform bit-flip diffusion* (Sun & Yang, 2023; Li et al., 2024):

$$q(\mathbf{x}_t | \mathbf{x}_{t-1}) = (1 - \beta_t) \mathbf{x}_{t-1} + \beta_t (1 - \mathbf{x}_{t-1}), \quad (2)$$

where $\beta_t \in (0, 1)$ denotes the corruption rate at step t . The corresponding t -step marginal is

$$q(\mathbf{x}_t | \mathbf{x}_0) = (1 - \bar{\beta}_t) \mathbf{x}_0 + \bar{\beta}_t (1 - \mathbf{x}_0), \quad (3)$$

where $\bar{\beta}_t = 1 - \prod_{s=1}^t (1 - \beta_s)$ is the cumulative corruption rate up to step t . Denoisers are then trained to recover the clean signal \mathbf{x}_0 or directly predict the optimal solution x^* using variants of likelihood or optimization-consistency losses (Sun & Yang, 2023; Li et al., 2024).

2.3 MASKED DIFFUSION

Recent progress has shown that *Masked Diffusion Models* (MDMs) often achieve better performance than uniform bit-flip diffusion in sequence inference tasks (Austin et al., 2021; Ou et al., 2025; Nie et al., 2025b). Instead of symmetrically flipping 0 and 1, MDMs corrupt data by replacing entries with a dedicated [MASK] token, following a continuous trajectory parameterized by $t \in [0, 1]$:

$$q(\mathbf{x}_t | \mathbf{x}_0) = (1 - t) \cdot \mathbf{x}_0 + t \cdot [\text{MASK}]. \quad (4)$$

162 In the context of combinatorial optimization, solutions are represented as binary vectors $\mathbf{x} \in \{0, 1\}^N$, where each entry (1) corresponds to a selected edge or node. Thus, any state where only
 163 a subset of 1s is visible can naturally be interpreted as a *partial solution*. MDMs exploit this per-
 164 spective: observed entries remain fixed while masked entries are left to be predicted by the denoiser.
 165 Moreover, since the corruption level is encoded in the fraction of masked variables rather than in
 166 an explicit timestep, training can be made *time-agnostic*. These properties have proven especially
 167 useful in large language models, enabling schedule-invariant training and efficient parallel decod-
 168 ing (Zheng et al., 2025; Ou et al., 2025; Nie et al., 2025a; Wu et al., 2025; Sun et al., 2025).
 169

170

171 3 METHOD

172

173 We introduce **NEXCO**, a masked diffusion framework tailored for combinatorial optimization. The
 174 core idea is to reinterpret the diffusion trajectory in line with the AE paradigm, as a constructive
 175 process over *partial solutions* with adaptive expansion, rather than as the generation of probability
 176 heatmaps. First, we design a *CO-specific corruption* (§3.1) that masks out selected variables but
 177 never introduces false positives, ensuring that intermediate states remain aligned with the feasible
 178 manifold. Then, we propose a *time-agnostic graph denoiser* (§3.2) trained with a new *optimization*
 179 *consistency* principle (§3.3), which enforces that all corrupted states of an instance consistently map
 180 to the same optimum. Finally, we develop a *native adaptive expansion* decoding strategy (§3.4),
 181 which progressively expands partial states into complete solutions under feasibility projection.

182

183 3.1 FORWARD PROCESS: CO-SPECIFIC CORRUPTION

184

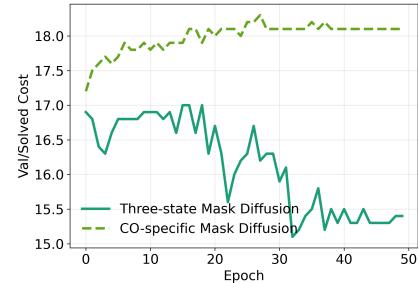
185 **Why uniform bit-flip fails to leverage intermediate states.** Prior GP-style diffusion solvers for
 186 CO (e.g., DIFUSCO (Sun & Yang, 2023), T2T (Li et al., 2023), Fast T2T (Li et al., 2024)) adopt
 187 *uniform bit-flip* corruption with the marginal formulation in Eq. 3, where each variable is flipped
 188 independently with probability $\bar{\beta}_t$. This design causes a fundamental *structural misalignment*: sym-
 189 metric flipping ignores the combinatorial constraints encoded in the graph. As a result, the corrupted
 190 state \mathbf{x}_t quickly drifts away from the feasible manifold. For example, in TSP it contains many edges
 191 that violate degree constraints or form subtours (see Fig.1 in (Li et al., 2024)). Such noisy states
 192 cannot be interpreted as valid *partial solutions*, but only as dense heatmaps detached from feasi-
 193 bility. Consequently, the denoiser is trained on spurious patterns, and the intermediate trajectory
 194 becomes unusable for constructive decoding, forcing prior solvers to discard it and rely solely on
 195 heuristic decoding at the final step. This highlights the need for a CO-aware corruption mechanism
 196 that *preserves sparsity and respects structural constraints*, so that intermediate states are meaningful
 197 partial solutions and can be directly exploited for adaptive expansion.

198

199 **Mask diffusion.** A natural alternative to uniform flipping
 200 is the *masked diffusion model (MDM)*(Austin et al., 2021;
 201 Nie et al., 2025b), which introduces a third state [MASK]
 202 and requires the denoiser to reconstruct both 0s and 1s.
 203 While effective in language tasks, directly applying MDM
 204 to CO is problematic. Take MIS as an example: the val-
 205 idation cost should increase during training, since the goal
 206 is to enlarge the independent set. However, as shown in
 207 Fig. 2, three-state MDM instead shows a *decreasing* val-
 208 idation cost. This mismatch stems from two factors. First,
 209 CO solutions are highly imbalanced: most variables are 0,
 210 so the training signal is dominated by negative examples,
 211 biasing the denoiser toward predicting 0s and shrinking
 212 the independent set. Second, unlike TSP where a global
 213 Hamiltonian cycle provides strong priors, MIS depends
 214 mainly on local adjacency constraints. Starting from a
 215 fully masked state, the model tends to favor “safe” pre-
 216 dictions (0s) over “risky” ones (1s that might violate inde-
 217 pendence), further reinforcing conservative behavior.

218

219 **CO-specific corruption.** To address this issue, we unify the [MASK] and 0 states into a single
 220 *background state* (numerically represented as 0), leaving only [BACKGROUND] \leftrightarrow 1 transitions.



221 **Figure 2: Validation cost curves on MIS.** CO-specific diffusion follows
 222 the correct trend of increasing set size, while three-state mask diffusion col-
 223 lapses to smaller sets. The latter suf-
 224 fers from severe 0–1 imbalance and lo-
 225 cal constraints, which bias the model
 226 toward conservative 0 predictions.

216 Intuitively, in CO most variables are fixed at 0 by problem-specific constraints such as degree limits
 217 in TSP, capacity restrictions in CVRP, or adjacency rules in MIS. Thus, treating [MASK] as distinct
 218 provides no additional signal but only exacerbates the imbalance by multiplying negative examples.
 219 Under this unification, the forward process becomes a one-way corruption in which only active
 220 entries (1s) may be dropped to [BACKGROUND]:

$$221 \quad q(\mathbf{x}_t | \mathbf{x}_0) = (1 - t) \cdot \mathbf{x}_0 + t \cdot \mathbf{0}, \quad (5)$$

223 where $\mathbf{0}$ denotes the background state. This eliminates the imbalance-driven conservatism of MDM
 224 and ensures that intermediate states correspond to valid *partial solutions* aligned with combinatorial
 225 feasibility. As shown in Fig. 2, our CO-specific corruption follows the correct trend of increasing
 226 validation cost, providing a clean foundation for the time-agnostic denoiser (§3.2) and the Native
 227 Adaptive Expansion decoding scheme (§3.4).

228 3.2 TIME-AGNOSTIC DENOISER

230 **From timestep to mask conditioning.** In prior CO diffusion solvers (Sun & Yang, 2023; Li et al.,
 231 2024), the denoiser $f_\theta(\mathbf{x}_t, t, G)$ explicitly conditions on timestep t , since t encodes the corruption
 232 intensity. This dependence rigidly ties the model to a predefined schedule, limiting generalization
 233 across horizons. In our CO-specific corruption process, however, the corruption level is directly *vis-
 234 ible in the mask pattern itself*: the fraction of surviving 1s naturally reflects the signal-to-noise ratio.
 235 We therefore remove timestep embeddings and design a *time-agnostic denoiser* $f_\theta(\mathbf{x}_t, G)$, which
 236 depends only on the corrupted state and the graph instance. This eliminates schedule sensitivity and
 237 shifts the focus to structural dependencies and the semantics of partial solutions.

238 **Model architecture.** We instantiate f_θ as an anisotropic graph neural network (GNN) (Joshi et al.,
 239 2021; Sun & Yang, 2023). Nodes and edges are annotated with task-specific features (e.g., coordinates
 240 in TSP, adjacency in MIS, or capacities in CVRP), while \mathbf{x}_t is encoded as binary attributes.
 241 Message passing aggregates both structural and partial-solution context, and attention-based pooling
 242 captures long-range dependencies. The output is a probability vector $\mathbf{p} \in [0, 1]^N$, where $p^{(i)}$
 243 estimates the likelihood of variable i belonging to the optimal solution x^* . Compared with conventional
 244 denoisers, the only change is the removal of timestep embeddings, highlighting that schedule
 245 awareness is unnecessary under mask corruption. Further details are provided in Appendix C.

246 **Analogy to large language diffusion models.** Our time-agnostic design parallels the success of
 247 masked diffusion in large language models (Austin et al., 2021; Ou et al., 2025; Nie et al., 2025b;
 248 Wu et al., 2025). In those settings, random masking without explicit timestep conditioning enables
 249 scalable pretraining and efficient parallel decoding. By extending this principle from token se-
 250 quences to graph-structured CO problems, we show that time-agnostic denoising is equally effective
 251 when intermediate states correspond to valid partial solutions.

252 3.3 TRAINING OBJECTIVE: TIME-AGNOSTIC OPTIMIZATION CONSISTENCY

254 **Consistency principle.** Consistency models (Song et al., 2023) learn direct mappings from noisy to
 255 clean data, enforcing that predictions across different corruption levels remain stable. Fast T2T (Li
 256 et al., 2024) adapted this idea to CO with *optimization consistency*, requiring all corrupted states of
 257 an instance to map to its optimal solution x^* . This ties denoising directly to the optimization goal,
 258 thereby improving one-step prediction quality.

259 **Time-agnostic optimization consistency.** Our CO-specific corruption is monotone: as t increases,
 260 supports shrink, so for any $0 < t < t' < 1$ we have $\text{supp}(\mathbf{x}_{t'}) \subseteq \text{supp}(\mathbf{x}_t)$. Each corrupted state
 261 \mathbf{x}_t is thus a valid subset of x^* . This property allows us to drop timestep embeddings and directly
 262 enforce consistency across corruption levels. Formally, the *time-agnostic optimization consistency*
 263 loss is

$$264 \quad \mathcal{L}_{\text{TOC}}(\theta) = \mathbb{E}_{t' > t} \left[d(f_\theta(\mathbf{x}_{t'}, G), x^*) + d(f_\theta(\mathbf{x}_t, G), x^*) \right], \quad (6)$$

266 where $d(\cdot, \cdot)$ is binary cross-entropy or a task-specific discrepancy. Training under \mathcal{L}_{TOC} amounts
 267 to reconstructing x^* from multiple *partial solutions*. Because the corruption never introduces false
 268 positives, every \mathbf{x}_t stays close to the feasible manifold, providing supervision that is inherently
 269 aligned with CO constraints. This contrasts with uniform diffusion, where intermediate states are
 often unrealistic and force the model to correct artifacts.

270 **Practical note.** The reference solution x^* can come from exact solvers on small/medium instances
 271 or high-quality heuristics on larger ones. The TOC loss remains valid in both cases, as it only
 272 requires a consistent reference per instance. We empirically confirm this in Table 3, where models
 273 trained on suboptimal labels still deliver competitive solutions.
 274

275 **3.4 INFERENCE PROCESS: NATIVE ADAPTIVE EXPANSION (NAE)**
 276

277 **Motivation.** Existing diffusion-based CO solvers generate a sequence of T_s global probability
 278 heatmaps $\mathbf{p}_1, \dots, \mathbf{p}_{T_s} \in [0, 1]^N$, but these intermediate states are not semantically valid partial
 279 solutions. As a result, prior methods typically exploit only the final step via heuristic decoding:
 280

$$\hat{x} = \Gamma(\text{Decode}(\mathbf{p}_{T_s})), \quad (7)$$

282 where `Decode` is a heuristic (e.g., greedy search) and Γ enforces feasibility (Sun & Yang, 2023).
 283 This *under-utilization* of the trajectory is a key limitation: despite producing many intermediate
 284 states, only a final one-shot prediction is retained. `COExpander` (Ma et al., 2025) alleviates this
 285 via *adaptive expansion* (AE), but AE is realized as an external wrapper around GP predictors. Its
 286 complexity depends on both the number of expansion rounds D_s and the per-call cost C_{GP} of the
 287 backbone: $\text{Cost}_{\text{AE}} = \mathcal{O}(D_s \cdot C_{\text{GP}})$. Thus, wrapping diffusion solvers (e.g., `Fast T2T` (Li et al.,
 288 2024)) yields $\mathcal{O}(D_s \cdot T_s)$ complexity, while wrapping GCN (Joshi et al., 2019) reduces it to $\mathcal{O}(D_s)$
 289 but still ties performance to external heatmap quality (Xia et al., 2024).

290 By contrast, our framework produces semantic partial solutions along the diffusion trajectory. This
 291 enables *Native Adaptive Expansion* (NAE): a deterministic expansion procedure that reuses the
 292 denoiser once per stage and enforces feasibility at each step, achieving $\mathcal{O}(T_s)$ complexity while making
 293 full constructive use of the entire trajectory.

294 **Procedure.** NAE begins from \mathbf{x}_1 as fully
 295 masked solution $\mathbf{0}$ and expands iteratively.
 296 At step t , the denoiser outputs a confidence
 297 vector $\mathbf{p}_t = f_\theta(G, \mathbf{x}_{t-1})$. Variables above
 298 threshold α form a candidate set \mathcal{C}_t , which
 299 is projected onto a feasible subset \mathcal{S}_t by
 300 $\Gamma(\cdot)$. From this subset, a fraction ρ_t of
 301 entries is selected, where ρ_t may be set as a
 302 tunable hyperparameter or determined by
 303 an evenly spaced schedule over steps. Re-
 304 peating this process yields a trajectory of
 305 feasible partial solutions until completion,
 306 as summarized in Algorithm 1.

307 **Algorithm 1: Native Adaptive Expansion (NAE)**
 308

309 **Input:** Graph G , denoiser f_θ , iterations T_s ,
 310 threshold α , expansion schedule $\{\rho_t\}$.
 311 Initialize $\mathbf{x}_1 \leftarrow \mathbf{0}$;
 312 **for** $t = 1, \dots, T_s$ **do**
 313 $\mathbf{p}_t \leftarrow f_\theta(G, \mathbf{x}_{t-1})$;
 314 $\mathcal{C}_t \leftarrow \{i \mid \mathbf{x}_{t-1}^{(i)} = 0, \mathbf{p}_t^{(i)} \geq \alpha\}$;
 315 Project candidates: $\mathcal{S}_t \leftarrow \Gamma(\mathcal{C}_t, \mathbf{p}_t, \mathbf{x}_{t-1})$;
 316 Retain ρ_t fraction of \mathcal{S}_t ;
 317 Update \mathbf{x}_t by activating retained entries;
 318 **return** \mathbf{x}_{T_s}

319 **Feasibility projection.** Although different CO tasks impose different feasibility rules, the projection
 320 operator $\Gamma(\cdot)$ follows a single task-agnostic template across all problems we study. At each ex-
 321 pansion step, the model produces a confidence vector \mathbf{p} , and $\Gamma(\cdot)$ constructs the next partial solution
 322 using the same three-stage procedure. First, candidate variables are sorted in descending confidence.
 323 Second, candidates are examined sequentially. Third, a candidate is activated only when doing so
 324 satisfies a simple boolean feasibility predicate. This local predicate is lightweight to compute and
 325 varies only in its constraint definition, not in the mechanism of feasibility projection. Consequently,
 326 extending NAE to a new CO task requires only defining this predicate, while the entire three-stage
 327 projection pipeline remains intact. Formally, the projection step solves:

$$\mathcal{S}_t = \arg \max_{x^{(i)} \subseteq \mathcal{C}_t, x \in \Omega} \sum_i p_t^{(i)} x^{(i)}, \quad (8)$$

328 which is implemented by inserting candidates in descending confidence order and accepting them
 329 only when the feasibility predicate holds. In practice, this yields a simple and uniform instantiation
 330 across tasks: TSP: edges are added while maintaining degree-2 and subtour-free constraints; MIS: a
 331 vertex is selected only if all neighbors remain inactive; CVRP: routing edges are inserted provided
 332 vehicle-capacity constraints are not violated. Candidates that fail the predicate are skipped, prevent-
 333 ing conflicts and ensuring that $\mathcal{S}_t \subseteq \mathcal{C}_t$ is always a feasible and high-confidence partial solution.

324 Table 2: Results on synthetic TSP problem instances. BS: Beam Search.
325

326 Algorithm	327 Type	328 TSP-100			329 TSP-500			330 TSP-1000		
		331 Length↓	332 Drop↓	333 Time	334 Length↓	335 Drop↓	336 Time	337 Length↓	338 Drop↓	339 Time
<i>Mathematical Solvers or Heuristics</i>										
Concorde (Applegate et al., 2006)	Exact	7.76*	—	0.23s	16.55*	—	18.65s	23.12*	—	84.38s
<i>Learning-based Solvers with Greedy Decoding</i>										
AM +BS (Kool et al., 2019)	LC	7.95	2.48%	0.61s	19.53	18.03%	1.31s	29.90	29.24%	5.91s
BQ-NCO +BS (Drakulic et al., 2023)	LC	7.76	0.01%	0.19s	16.64	0.55%	7.03s	23.47	1.38%	17.81s
GCN +BS (Joshi et al., 2019)	GP	8.41	8.38%	0.28s	30.37	83.55%	17.81s	51.26	121.73%	24.23s
DIMES (Qiu et al., 2022)	GP	8.01	3.23%	0.06s	17.17	3.74%	0.45s	24.79	7.22%	1.12s
DIFUSCO ($T_s = 50$) (Sun & Yang, 2023)	GP	7.78	0.26%	0.59s	16.82	1.64%	1.43s	23.57	1.94%	5.04s
T2T ($T_s = 50, T_g = 30$) (Li et al., 2023)	GP	7.76	0.07%	1.34s	16.68	0.82%	3.05s	23.44	1.40%	9.23s
Fast T2T ($T_s = 5$) (Li et al., 2024)	GP	7.76	0.08%	0.06s	16.72	1.02%	0.27s	23.38	1.12%	0.99s
Fast T2T ($T_s = 5, T_g = 5$) (Li et al., 2024)	GP	7.76	0.03%	0.31s	16.61	0.39%	1.41s	23.25	0.58%	5.81s
COExpander ($D_s = 3, T_s = 5$) (Ma et al., 2025)	AE	7.76	0.04%	0.18s	16.63	0.52%	0.61s	23.34	0.95%	2.26s
NEXCO ($D_s = 3$)	NAE	7.76	0.04%	0.05s	16.61	0.39%	0.23s	23.31	0.85%	0.91s
NEXCO ($D_s = 5$)	NAE	7.76	0.03%	0.08s	16.59	0.28%	0.33s	23.26	0.63%	1.31s
NEXCO ($D_s = 7$)	NAE	7.76	0.02%	0.11s	16.59	0.25%	0.43s	23.24	0.52%	1.68s
<i>Learning-based Solvers with 4× Sampling Decoding</i>										
LEHD PRC 100 (Luo et al., 2023)	LC	7.76	0.01%	0.64s	16.61	0.34%	3.75s	23.44	1.22%	20.16s
Fast T2T ($T_s = 5, T_g = 5$) (Li et al., 2024)	GP	7.76	0.01%	0.99s	16.58	0.21%	5.16s	23.22	0.42%	17.42s
COExpander ($D_s = 3, T_s = 5$) (Ma et al., 2025)	AE	7.76	0.01%	0.61s	16.59	0.24%	2.21s	23.27	0.64%	8.43s
NEXCO ($D_s = 5$)	NAE	7.76	0.01%	0.25s	16.57	0.14%	1.16s	23.20	0.35%	4.85s

341 **Convergence and complexity.** We provide a formal convergence analysis of the NAE procedure.
 342 Because the forward corruption in our CO-specific diffusion process is one-way absorbing ($1 \rightarrow 0$),
 343 every diffusion state remains a feasible partial solution. The reverse step expands this partial solution
 344 via $\mathbf{x}_{t+1} = \Gamma(\mathbf{x}_t \vee \mathbf{z}_t)$, where \mathbf{z}_t is the candidate activation mask predicted by the denoiser, and Γ
 345 is a feasibility projector. We assume the following mild and standard conditions, satisfied by the
 346 projectors $\Gamma(\cdot)$ used for TSP, MIS, and CVRP:

347 **(A1) Monotone projection:** $\Gamma(\mathbf{x}) \succeq \mathbf{x}$ for all feasible \mathbf{x} .

348 **(A2) Strict expandability:** $\exists \mathbf{z}_t$ such that $\Gamma(\mathbf{x}_t \vee \mathbf{z}_t) \succ \mathbf{x}_t$ whenever \mathbf{x}_t is incomplete.

349 **(A3) Bounded solution size:** Any complete feasible solution contains at most N_{\max} active variables.

353 **Proposition 1 (Finite-time convergence of NAE)** *Under assumptions (A1)-(A3), NAE generates a*
 354 *monotone sequence $\mathbf{x}_0 \preceq \mathbf{x}_1 \preceq \dots$ and converges to a complete feasible solution in at most N_{\max}*
 355 *iterations.*

356 **Remark.** The upper bound N_{\max} is fully consistent with typical CO structures:

- 358 • TSP: $N_{\max} = N$ edges in a Hamiltonian tour.
- 359 • MIS: $N_{\max} \leq n$ selected nodes.
- 360 • CVRP: N_{\max} equals the total number of edges across all valid routes.

363 This analysis formalizes the intuition that NAE is a monotone constructive decoder that reaches a
 364 complete feasible solution in finite time. In terms of efficiency, NAE requires $\mathcal{O}(T_s)$ denoiser calls,
 365 matching the order of diffusion while being asymptotically more efficient than COExpander, whose
 366 wrapper design incurs $\mathcal{O}(D_s \cdot T_s)$ complexity.

368 4 EXPERIMENTS

369 4.1 EXPERIMENTS ON TSP

372 **Datasets.** Each TSP instance consists of N two-dimensional coordinates and a reference optimal
 373 solution. Following standard practice (Sun & Yang, 2023), we generate instances by uniformly
 374 sampling N nodes from the unit square $[0, 1]^2$. The training sets contain 1,280K, 128K, and 64K
 375 instances for TSP-100, TSP-500, and TSP-1000, respectively. The corresponding test sets consist
 376 of 1,280 instances for TSP-100 and 128 instances each for TSP-500 and TSP-1000. Reference
 377 solutions are obtained using Concorde (Applegate et al., 2006). We further include large-scale TSP-
 10K instances and real-world TSPLIB benchmarks, with results reported in Appendix B.2 and B.3.

378 **Metrics.** We evaluate solvers on three metrics: 1) *Length*: average tour length of the produced
 379 solutions; 2) *Drop*: relative deviation from the reference solution; 3) *Time*: average time per instance.
 380

381 **Setting.** For NEXCO, we vary the number of expansion rounds $D_s \in \{3, 5, 7\}$ in the NAE procedure
 382 to balance solution quality and runtime. The confidence threshold α is tuned as a hyperparameter.
 383 For diffusion-based baselines, we adopt standard configurations (Li et al., 2023): T_s denotes the
 384 number of inference steps, and T_g denotes the number of gradient refinement steps. Unless otherwise
 385 stated, all methods employ greedy decoding with an optional 2-Opt heuristic in post-processing.
 386

387 **Main results.** Table 2 reports the comparison across different scales. NEXCO consistently surpasses
 388 state-of-the-art learning-based solvers in both solution quality and runtime. On TSP-100, Fast T2T
 389 achieves a 0.03% gap in 0.31s, whereas NEXCO attains the same gap in only 0.08s, yielding a $3.9 \times$
 390 speedup. Similarly, on TSP-500, NEXCO reduces the gap to 0.25% within 0.43s, while Fast T2T
 391 requires 1.41s to reach 0.39%, corresponding to both a $1.5 \times$ improvement in gap and a $3 \times$ speedup.
 392 Taken together, these results demonstrate that embedding adaptive expansion natively into diffusion
 393 not only accelerates inference by $2\text{--}4 \times$, but also consistently reduces optimality gaps, highlighting
 394 the effectiveness of NEXCO as a next-generation neural TSP solver.
 395

396 **Ablation on reference quality.** We further ex-
 397 amine the robustness of NEXCO when trained
 398 with suboptimal supervision. To this end, we
 399 construct perturbed labels by applying 2-Opt
 400 local search to the ground-truth solutions. As
 401 shown in Table 3, these perturbed references are
 402 significantly worse than the exact optima, with gaps of 1.65% and 3.35% on TSP-500. Neverthe-
 403 less, models trained on such labels still deliver highly competitive results, achieving final gaps of
 404 only 0.25% – 0.31%. This demonstrates that NEXCO is not tied to exact optimal labels but can
 405 effectively leverage high-quality heuristic solutions as consistent training signals, making it broadly
 406 applicable in practical CO scenarios where exact optima are often unavailable.
 407

408 **Generalization study.** Table 4 reports cross-
 409 scale transfer results. Models trained on small
 410 instances generalize poorly, with DIFUSCO,
 411 T2T, and Fast T2T all incurring $\sim 3\%$ gaps
 412 on TSP-1000. Training on medium or large
 413 instances improves robustness, and NEXCO
 414 shows the strongest transferability. In partic-
 415 ular, when trained on TSP-500, it generalizes
 416 to TSP-1000 with only 0.58% gaps. Simi-
 417 larly, the model trained on TSP-1000 achieves
 418 a 0.37% gap on TSP-500, surpassing baselines
 419 that are directly trained on TSP-500. These
 420 results highlight that native adaptive expansion
 421 yields consistently superior cross-scale gen-
 422 eralization compared with existing diffusion-based solvers.
 423

4.2 EXPERIMENTS ON MIS.

424 **Datasets.** Following (Sun & Yang, 2023; Qiu et al., 2022; Li et al., 2018; Ahn et al., 2020), we
 425 consider two benchmark datasets: RB graphs and Erdős-Rényi (ER) graphs. For RB graphs, we
 426 generate small- and large-scale instances by sampling the number of vertices from [200, 300] and
 427 [800, 1200], respectively. For ER graphs, we construct random graphs with 700–800 nodes, where
 428 each edge is added independently with probability 0.15. In total, we generate 90,000 small and
 429 6,400 large RB instances for training, with 500 test instances. For ER graphs, 163,840 instances are
 430 used for training and 500 test instances are adopted from (Qiu et al., 2022).
 431

432 **Metrics.** 1) *Size*: the average solution size; 2) *Drop*: the relative deviation from the reference
 433 solution obtained by KaMIS (Lamm et al., 2016); 3) *Time*: the average runtime per instance.
 434

435 **Main results.** Table 5 shows that NEXCO consistently improves both efficiency and solution quality
 436 over diffusion- and expansion-based baselines. On RB-[200–300], NEXCO lowers the drop from
 437

Table 3: Ablation on reference quality.

Label / Model	Length \downarrow	Drop \downarrow	Length \downarrow	Drop \downarrow
2-Opt Perturbation	16.82	1.65%	17.11	3.35%
NEXCO ($D_s = 5$)	16.60	0.31%	16.60	0.31%
NEXCO ($D_s = 7$)	16.59	0.25%	16.59	0.26%

Table 4: Cross-scale generalization results on TSP. Each entry reports Length, Gap (%).

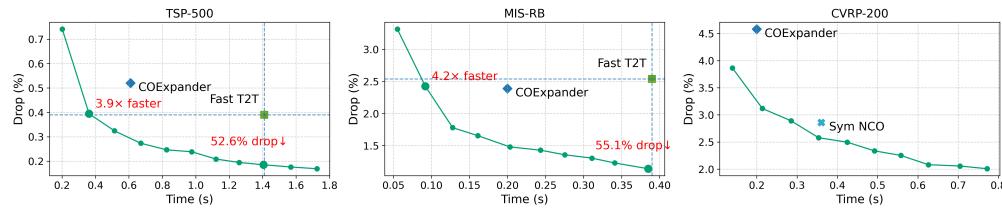
Training \ Testing	TSP-100	TSP-500	TSP-1000
DIFUSCO ($T_s=50$)	7.78, 0.23%	7.85, 1.16%	7.87, 1.42%
T2T ($T_s=50, T_g=30$)	7.77, 0.08%	7.95, 2.47%	7.91, 1.96%
Fast T2T ($T_s=5, T_g=5$)	7.77, 0.02%	7.79, 0.40%	7.80, 0.55%
NEXCO ($D_s = 5$)	7.76, 0.02%	7.77, 0.18%	7.81, 0.72%
DIFUSCO ($T_s=50$)	17.05, 3.04%	16.78, 1.40%	16.86, 1.85%
T2T ($T_s=50, T_g=30$)	16.92, 2.25%	16.68, 0.81%	16.72, 1.00%
Fast T2T ($T_s=5, T_g=5$)	16.92, 1.77%	16.61, 0.38%	16.63, 0.51%
NEXCO ($D_s = 5$)	17.04, 2.96%	16.59, 0.27%	16.60, 0.37%
DIFUSCO ($T_s=50$)	24.04, 3.98%	23.65, 2.30%	23.63, 2.21%
T2T ($T_s=50, T_g=30$)	23.85, 3.16%	23.47, 1.51%	23.41, 1.23%
Fast T2T ($T_s=5, T_g=5$)	23.77, 3.08%	23.31, 0.81%	23.25, 0.58%
NEXCO ($D_s = 5$)	24.01, 3.86%	23.25, 0.58%	23.25, 0.57%

432 Table 5: Results on synthetic MIS problem instances.
433

434 Algorithm	435 Type	436 RB-[200-300]			437 RB-[800-1200]			438 ER-[700-800]		
		439 Size↑	440 Drop↓	441 Time	442 Size↑	443 Drop↓	444 Time	445 Size↑	446 Drop↓	447 Time
KaMIS (Lamm et al., 2016)	Heuristics	20.10	—	45.81s	43.00	—	56.97s	44.87	—	48.49s
DIMES (Qiu et al., 2022)	GP	—	—	—	—	—	—	38.24	14.78%	2.8s
DIFUSCO ($T_s = 100$) (Sun & Yang, 2023)	GP	18.52	7.81%	1.9s	—	—	—	37.03	18.53	2.58s
T2T ($T_s = 50, T_g = 30$) (Li et al., 2023)	GP	18.98	5.49%	2.52s	—	—	—	39.81	11.28	3.33s
Fast T2T ($T_s = 5, T_g = 5$) (Li et al., 2024)	GP	19.58	2.54%	0.39s	39.34	8.51%	2.76s	40.78	9.31%	1.22s
COExpander ($D_s = 20, T_s = 1$) (Ma et al., 2025)	AE	19.60	2.39%	0.20s	41.09	4.36%	2.05s	42.44	5.62%	1.53s
NEXCO ($D_s = 5$)	NAE	19.70	1.97%	0.11s	40.94	4.79%	0.75s	42.64	4.96%	0.44s
NEXCO ($D_s = 7$)	NAE	19.76	1.66%	0.14s	41.25	4.07%	1.00s	42.98	4.20%	0.56s

440 Table 6: Results on synthetic CVRP problem instances.
441

442 Algorithm	443 Type	444 CVRP 50			445 CVRP 100			446 CVRP 200		
		447 Length↓	448 Drop↓	449 Time	450 Length↓	451 Drop↓	452 Time	453 Length↓	454 Drop↓	455 Time
HGS (Helsgaun, 2017)	Heuristics	10.37	—	1.01s	15.56	—	20.03s	19.63	—	60.02s
Sym-NCO (Kim et al., 2022)	LC	10.57	1.95%	0.09s	15.93	2.37%	0.19s	20.19	2.86%	0.36s
COExpander ($D_s = 3, T_s = 1$) (Ma et al., 2025)	AE	10.77	3.85%	0.03s	16.22	4.19%	0.06s	20.52	4.58%	0.20s
NEXCO ($D_s=3$)	NAE	10.48	1.12%	0.04s	15.83	1.73%	0.08s	20.18	2.76%	0.27s
NEXCO ($D_s=5$)	NAE	10.46	0.85%	0.06s	15.78	1.40%	0.12s	20.11	2.45%	0.39s

456 Figure 3: **Effect of expansion steps on performance.** Green curves show our method with different
457 numbers of expansion steps. Increasing steps improves solution quality (smaller drop) while linearly
458 increasing inference time. Reference markers denote SOTA baselines (Fast T2T, COExpander, Sym-
459 NCO), against which our method achieves consistently better efficiency–quality trade-offs.
460

2.54% (Fast T2T) to 1.66% while being 2.8 \times faster. On ER, it achieves a 4.20% drop compared with 9.31% for Fast T2T, with more than 2 \times speedup. Compared with COExpander, NEXCO also attains smaller gaps in much less time: on RB-[800–1200], 4.07% in 1.00s vs. 4.36% in 2.05s. These results demonstrate that NAE not only improves scalability but also yields consistently higher-quality solutions in MIS. Additional generalization results are provided in Appendix B.1.

4.3 EXPERIMENTS ON CVRP

Datasets. Each CVRP instance consists of a depot, N customer coordinates, and corresponding demands. Coordinates are uniformly sampled from the unit square $[0, 1]^2$, demands are drawn as integers from $[1, 10]$, and vehicle capacities are fixed to 40, 50, and 80 for CVRP-50, CVRP-100, and CVRP-200, respectively. The training sets contain 1,280K, 640K, and 32K instances for CVRP-50/100/200, while the test sets include 10K instances for CVRP-50/100 and 100 instances for CVRP-200. Reference solutions are obtained using the HGS solver (Helsgaun, 2017).

Main results. As summarized in Table 6, NEXCO achieves consistent improvements across all scales. Compared with LC methods such as Sym-NCO (Kim et al., 2022), which yield moderate drops, and AE-based COExpander, which suffers from even larger deviations, NEXCO substantially reduces the gap while also cutting runtime. For example, it improves from 3.86% \rightarrow 0.85% on CVRP-50, 4.16% \rightarrow 1.40% on CVRP-100, and 4.84% \rightarrow 2.45% on CVRP-200, corresponding to 1.6 – 2 \times speedup. Notably, GP-based diffusion solvers have not been extended to CVRP due to its intricate capacity constraints, further highlighting the strength of native adaptive expansion.

4.4 HYPERPARAMETER STUDY

We analyze two key hyperparameters in NEXCO: the number of expansion steps and the candidate threshold α . As shown in Fig. 3, increasing the number of expansion steps consistently improves solution quality while runtime grows nearly linearly, revealing a clear trade-off between efficiency and

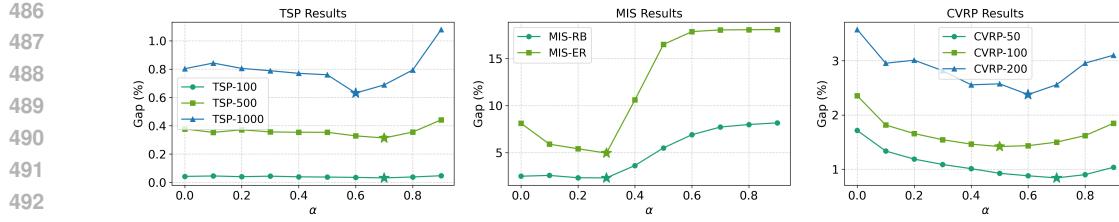


Figure 4: **Effect of the candidate threshold α .** Performance across TSP, MIS, and CVRP improves at moderate thresholds, but deteriorates when α is set too high (restrictive) or too low (noisy).

performance. This allows practitioners to flexibly adjust inference cost, and NEXCO remains superior to strong baselines (Fast T2T, COExpander, SymNCO) even under reduced step budgets. Similarly, Fig. 4 shows that moderate values of α yield the best results: overly high thresholds become too restrictive, while overly low thresholds admit noisy candidates. Importantly, the optimal range of α is largely shared across all sizes and distributions within each task (e.g., TSP-100/500/1000; MIS-RB/ER; CVRP-50/100/200), demonstrating strong task-level robustness. This cross-setting consistency means that NEXCO requires minimal per-instance tuning, as a single hyperparameter choice generalizes reliably across scales within the same CO task.

5 CONCLUSION

In this work, we proposed **NEXCO**, a masked diffusion framework that realizes adaptive solution expansion as a native generative principle for neural combinatorial optimization. Our framework is built on three key components: a CO-specific forward corruption that preserves sparsity and yields semantic partial solutions, a time-agnostic GNN denoiser trained under optimization consistency, and a Native Adaptive Expansion (NAE) inference strategy that progressively selects confident variables under feasibility constraints. This framework has demonstrated the effectiveness across three representative CO problems. We believe this work opens up new opportunities for integrating constructive expansion mechanisms into diffusion-based generative modeling, and provides a step forward toward scalable and general-purpose neural solvers for combinatorial optimization.

ETHICS STATEMENT

This paper presents a new masked diffusion framework for neural combinatorial optimization. The proposed method addresses fundamental challenges in existing neural solvers, including the inefficiency of local construction, the constraint conflicts of global prediction, and the reliance on external predictors in adaptive expansion. Our contribution is methodological in nature, aiming to improve both solution quality and inference efficiency across benchmark CO problems such as MIS, TSP, and CVRP. We expect this work to benefit the broader research community by providing a more principled foundation for scalable and effective neural solvers in discrete optimization.

We do not anticipate any negative societal impacts arising from this research. Our work does not involve sensitive personal data, human subjects, or applications with immediate ethical risks. Furthermore, we do not identify issues related to conflicts of interest, discrimination or fairness, privacy or security, legal compliance, or research integrity. As a methodological contribution validated on public benchmarks, this work is aligned with the ethical standards of the ML community.

REPRODUCIBILITY STATEMENT

We provide detailed descriptions of the datasets, evaluation metrics, and hyperparameter choices in Sec. 4. The model architecture, hardware configuration, and complete training details are provided in Appendix C and Appendix D. All source code, preprocessed datasets, and pretrained checkpoints will be publicly released upon publication to ensure full reproducibility.

540 REFERENCES
541

542 Sungsoo Ahn, Younggyo Seo, and Jinwoo Shin. Learning what to defer for maximum independent
543 sets. In *International Conference on Machine Learning*, pp. 134–144. PMLR, 2020.

544 David Applegate, Ribert Bixby, Vasek Chvatal, and William Cook. Concorde tsp solver. <https://www.math.uwaterloo.ca/tsp/concorde/index.html>, 2006. Accessed: 2025-
545 04-21.

546 Jacob Austin, Daniel D Johnson, Jonathan Ho, Daniel Tarlow, and Rianne Van Den Berg. Structured
547 denoising diffusion models in discrete state-spaces. *Advances in neural information processing
548 systems*, 34:17981–17993, 2021.

549 Yoshua Bengio, Andrea Lodi, and Antoine Prouvost. Machine learning for combinatorial opti-
550 mization: a methodological tour d’horizon. *European Journal of Operational Research*, 290(2):
551 405–421, 2021.

552 Federico Berto, Chuanbo Hua, Junyoung Park, Laurin Luttmann, Yining Ma, Fanchen Bu, Jiarui
553 Wang, Haoran Ye, Minsu Kim, Sanghyeok Choi, et al. Rl4co: an extensive reinforcement learning
554 for combinatorial optimization benchmark. *arXiv preprint arXiv:2306.17100*, 2023.

555 Xinyun Chen and Yuandong Tian. Learning to perform local rewriting for combinatorial optimiza-
556 tion. *Advances in neural information processing systems*, 32, 2019.

557 Darko Drakulic, Sofia Michel, Florian Mai, Arnaud Sors, and Jean-Marc Andreoli. Bq-nco: Bisim-
558 ulation quotienting for efficient neural combinatorial optimization. In *Advances in Neural Infor-
559 mation Processing Systems*, volume 36, pp. 77416–77429, 2023.

560 Zhang-Hua Fu, Kai-Bin Qiu, and Hongyuan Zha. Generalize a small pre-trained model to arbitrarily
561 large tsp instances. In *Proceedings of the AAAI conference on artificial intelligence*, volume 35,
562 pp. 7474–7482, 2021.

563 Shansan Gong, Shivam Agarwal, Yizhe Zhang, Jiacheng Ye, Lin Zheng, Mukai Li, Chenxin An,
564 Peilin Zhao, Wei Bi, Jiawei Han, et al. Scaling diffusion language models via adaptation from
565 autoregressive models. In *The Thirteenth International Conference on Learning Representations*,
566 2025.

567 Keld Helsgaun. An extension of the lin-kernighan-helsgaun tsp solver for constrained traveling
568 salesman and vehicle routing problems. *Roskilde: Roskilde University*, 12:966–980, 2017.

569 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in
570 neural information processing systems*, 33:6840–6851, 2020.

571 Emiel Hoogeboom, Didrik Nielsen, Priyank Jaini, Patrick Forré, and Max Welling. Argmax flows
572 and multinomial diffusion: Learning categorical distributions. *Advances in neural information
573 processing systems*, 34:12454–12465, 2021.

574 André Hottung, Bhanu Bhandari, and Kevin Tierney. Learning a latent search space for routing prob-
575 lems using variational autoencoders. In *International Conference on Learning Representations*,
576 2021.

577 Qingchun Hou, Jingwei Yang, Yiqiang Su, Xiaoqing Wang, and Yuming Deng. Generalize learned
578 heuristics to solve large-scale vehicle routing problems in real-time. In *The Eleventh International
579 Conference on Learning Representations*, 2023.

580 Chaitanya K Joshi, Thomas Laurent, and Xavier Bresson. An efficient graph convolutional network
581 technique for the travelling salesman problem. *arXiv preprint arXiv:1906.01227*, 2019.

582 Chaitanya K Joshi, Quentin Cappart, Louis-Martin Rousseau, and Thomas Laurent. Learning tsp
583 requires rethinking generalization. In *27th International Conference on Principles and Practice
584 of Constraint Programming (CP 2021)*, pp. 33–1. Schloss Dagstuhl–Leibniz-Zentrum für Infor-
585 matik, 2021.

594 Minsu Kim, Jinkyoo Park, et al. Learning collaborative policies to solve np-hard routing problems.
 595 *Advances in Neural Information Processing Systems*, 34:10418–10430, 2021.
 596

597 Minsu Kim, Junyoung Park, and Jinkyoo Park. Sym-nco: Leveraging symmetricity for neural com-
 598 binatorial optimization. In *Advances in Neural Information Processing Systems*, volume 35, pp.
 599 1936–1949, 2022.

600 Wouter Kool, Herke van Hoof, and Max Welling. Attention, learn to solve routing problems! In
 601 *International Conference on Learning Representations*, 2019.
 602

603 Yeong-Dae Kwon, Jinho Choo, Byoungjip Kim, Iljoo Yoon, Youngjune Gwon, and Seungjai Min.
 604 Pomo: Policy optimization with multiple optima for reinforcement learning. In *Advances in*
 605 *Neural Information Processing Systems*, volume 33, pp. 21188–21198, 2020.

606 Sebastian Lamm, Peter Sanders, Christian Schulz, Darren Strash, and Renato F Werneck. Find-
 607 ing near-optimal independent sets at scale. In *2016 Proceedings of the eighteenth workshop on*
 608 *algorithm engineering and experiments (ALENEX)*, pp. 138–150. SIAM, 2016.

609

610 Sirui Li, Zhongxia Yan, and Cathy Wu. Learning to delegate for large-scale vehicle routing. *Ad-*
 611 *vances in Neural Information Processing Systems*, 34:26198–26211, 2021.

612 Yang Li, Jinpei Guo, Runzhong Wang, and Junchi Yan. T2t: From distribution learning in training
 613 to gradient search in testing for combinatorial optimization. In *Advances in Neural Information*
 614 *Processing Systems*, volume 36, pp. 50020–50040, 2023.

615

616 Yang Li, Jinpei Guo, Runzhong Wang, Hongyuan Zha, and Junchi Yan. Fast t2t: Optimization
 617 consistency speeds up diffusion-based training-to-testing solving for combinatorial optimization.
 618 In *Advances in Neural Information Processing Systems*, volume 37, pp. 30179–30206, 2024.

619

620 Yang Li, Lvda Chen, Haonan Wang, Runzhong Wang, and Junchi Yan. Generation as search operator
 621 for test-time scaling of diffusion-based combinatorial optimization. In *The Thirty-ninth Annual*
 622 *Conference on Neural Information Processing Systems*, 2025.

623

624 Zhuwen Li, Qifeng Chen, and Vladlen Koltun. Combinatorial optimization with graph convolu-
 625 tional networks and guided tree search. In *Advances in neural information processing systems*,
 volume 31, 2018.

626

627 Aaron Lou, Chenlin Meng, and Stefano Ermon. Discrete diffusion modeling by estimating the
 628 ratios of the data distribution. In *Proceedings of the 41st International Conference on Machine*
 629 *Learning*, pp. 32819–32848, 2024.

630

631 Fu Luo, Xi Lin, Fei Liu, Qingfu Zhang, and Zhenkun Wang. Neural combinatorial optimization with
 632 heavy decoder: Toward large scale generalization. *Advances in Neural Information Processing*
 633 *Systems*, 36:8845–8864, 2023.

634

635 Fu Luo, Xi Lin, Yaxin Wu, Zhenkun Wang, Tong Xialiang, Mingxuan Yuan, and Qingfu Zhang.
 636 Boosting neural combinatorial optimization for large-scale vehicle routing problems. In *The Thir-
 637 teenth International Conference on Learning Representations*, 2025.

638

639 Jiale Ma, Wenzheng Pan, Yang Li, and Junchi Yan. Coexpander: Adaptive solution expansion for
 640 combinatorial optimization. In *International Conference on Machine Learning*. PMLR, 2025.

641

642 Yining Ma, Zhiguang Cao, and Yeow Meng Chee. Learning to search feasible and infeasible re-
 643 gions of routing problems with flexible neural k-opt. *Advances in Neural Information Processing*
 644 *Systems*, 36:49555–49578, 2023.

645

646 Yimeng Min, Yiwei Bai, and Carla P Gomes. Unsupervised learning for solving the travelling sales-
 647 man problem. In *Advances in Neural Information Processing Systems*, volume 36, pp. 47264–
 648 47278, 2023.

649

650 Shen Nie, Fengqi Zhu, Chao Du, Tianyu Pang, Qian Liu, Guangtao Zeng, Min Lin, and Chongx-
 651 uan Li. Scaling up masked diffusion models on text. In *International Conference on Learning*
 652 *Representations*, 2025a.

648 Shen Nie, Fengqi Zhu, Zebin You, Xiaolu Zhang, Jingyang Ou, Jun Hu, Jun Zhou, Yankai
 649 Lin, Ji-Rong Wen, and Chongxuan Li. Large language diffusion models. *arXiv preprint*
 650 *arXiv:2502.09992*, 2025b.

651

652 Paulo Roberto Oliveira da Costa, Jason Rhuggenaath, Yingqian Zhang, and Alp Akcay. Learning
 653 2-opt heuristics for the traveling salesman problem via deep reinforcement learning. In *Asian*
 654 *conference on machine learning*, pp. 465–480. PMLR, 2020.

655 Jingyang Ou, Shen Nie, Kaiwen Xue, Fengqi Zhu, Jiacheng Sun, Zhenguo Li, and Chongxuan Li.
 656 Your absorbing discrete diffusion secretly models the conditional distributions of clean data. In
 657 *International Conference on Learning Representations*, 2025.

658

659 Wenzheng Pan, Hao Xiong, Jiale Ma, Wentao Zhao, Yang Li, and Junchi Yan. Unico: On unified
 660 combinatorial optimization via problem reduction to matrix-encoded general tsp. In *International*
 661 *Conference on Learning Representations*, 2025.

662 Ruizhong Qiu, Zhiqing Sun, and Yiming Yang. Dimes: A differentiable meta solver for combinatorial
 663 optimization problems. In *Advances in Neural Information Processing Systems*, volume 35,
 664 pp. 25531–25546, 2022.

665

666 Gerhard Reinelt. Tsplib—a traveling salesman problem library. *ORSA journal on computing*, 3(4):
 667 376–384, 1991.

668

669 Subham Sahoo, Marianne Arriola, Yair Schiff, Aaron Gokaslan, Edgar Marroquin, Justin Chiu,
 670 Alexander Rush, and Volodymyr Kuleshov. Simple and effective masked diffusion language
 671 models. *Advances in Neural Information Processing Systems*, 37:130136–130184, 2024.

672

673 Sebastian Sanokowski, Sepp Hochreiter, and Sebastian Lehner. A diffusion model framework for
 674 unsupervised neural combinatorial optimization. In *International Conference on Machine Learning*,
 675 pp. 43346–43367, 2024.

676

677 Martin JA Schuetz, J Kyle Brubaker, and Helmut G Katzgraber. Combinatorial optimization with
 678 physics-inspired graph neural networks. *Nature Machine Intelligence*, 4(4):367–377, 2022.

679

680 Jiaxin Shi, Kehang Han, Zhe Wang, Arnaud Doucet, and Michalis Titsias. Simplified and generalized
 681 masked diffusion for discrete data. *Advances in neural information processing systems*, 37:
 682 103131–103167, 2024.

683

684 Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised
 685 learning using nonequilibrium thermodynamics. In *International conference on machine learning*,
 686 pp. 2256–2265. PMLR, 2015.

687

688 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 689 Poole. Score-based generative modeling through stochastic differential equations. In *International*
 690 *Conference on Learning Representations*, 2021.

691

692 Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. In *International*
 693 *Conference on Machine Learning*, pp. 32211–32252. PMLR, 2023.

694

695 Jingyan Sui, Shizhe Ding, Ruizhi Liu, Liming Xu, and Dongbo Bu. Learning 3-opt heuristics for
 696 traveling salesman problem via deep reinforcement learning. In *Asian conference on machine*
 697 *learning*, pp. 1301–1316. PMLR, 2021.

698

699 Qiao Sun, Zhicheng Jiang, Hanhong Zhao, and Kaiming He. Is noise conditioning necessary for
 700 denoising generative models? In *International Conference on Machine Learning*. PMLR, 2025.

701

Zhiqing Sun and Yiming Yang. Difusco: Graph-based diffusion solvers for combinatorial optimization. In *Advances in neural information processing systems*, volume 36, pp. 3706–3731, 2023.

Thibaut Vidal, Teodor Gabriel Crainic, Michel Gendreau, Nadia Lahrichi, and Walter Rei. A hybrid
 702 genetic algorithm for multidepot and periodic vehicle routing problems. *Operations Research*, 60
 703 (3):611–624, 2012.

702 Chengyue Wu, Hao Zhang, Shuchen Xue, Zhijian Liu, Shizhe Diao, Ligeng Zhu, Ping Luo, Song
 703 Han, and Enze Xie. Fast-dllm: Training-free acceleration of diffusion llm by enabling kv cache
 704 and parallel decoding. *arXiv preprint arXiv:2505.22618*, 2025.

705

706 Yaxin Wu, Wen Song, Zhiguang Cao, Jie Zhang, and Andrew Lim. Learning improvement heuris-
 707 tics for solving routing problems. *IEEE transactions on neural networks and learning systems*,
 708 33(9):5057–5069, 2021.

709

710 Yifan Xia, Xianliang Yang, Zichuan Liu, Zhihao Liu, Lei Song, and Jiang Bian. Position: rethinking
 711 post-hoc search-based neural approaches for solving large-scale traveling salesman problems. In
 712 *Proceedings of the 41st International Conference on Machine Learning*, 2024.

713

714 Yubin Xiao, Di Wang, Boyang Li, Mingzhao Wang, Xuan Wu, Changliang Zhou, and You Zhou.
 715 Distilling autoregressive models to obtain high-performance non-autoregressive solvers for ve-
 716 hicle routing problems with faster inference speed. In *Proceedings of the AAAI conference on*
 717 *artificial intelligence*, volume 38, pp. 20274–20283, 2024.

718

719 Haoran Ye, Jiarui Wang, Helan Liang, Zhiguang Cao, Yong Li, and Fanzhang Li. Glop: Learning
 720 global partition and local construction for solving large-scale routing problems in real-time. In
 721 *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pp. 20284–20292, 2024.

722

723 Jiacheng Ye, Jiahui Gao, Shansan Gong, Lin Zheng, Xin Jiang, Zhenguo Li, and Lingpeng Kong.
 724 Beyond autoregression: Discrete diffusion for complex reasoning and planning. In *The Thirteenth*
 725 *International Conference on Learning Representations*, 2025.

726

727 Hang Zhao, Kexiong Yu, Yuhang Huang, Renjiao Yi, Chenyang Zhu, and Kai Xu. Disco: Ef-
 728 ficient diffusion solver for large-scale combinatorial optimization problems. *arXiv preprint*
 729 *arXiv:2406.19705*, 2024.

730

731 Kaiwen Zheng, Yongxin Chen, Hanzi Mao, Ming-Yu Liu, Jun Zhu, and Qinsheng Zhang. Masked
 732 diffusion models are secretly time-agnostic masked models and exploit inaccurate categorical
 733 sampling. In *International Conference on Learning Representations*, 2025.

734

735

736

737

738

739

740

741

742

743

744

745

746

747

748

749

750

751

752

753

754

755

APPENDIX

A ADDITIONAL RELATED WORKS

A.1 LEARNING-BASED COMBINATORIAL OPTIMIZATION

Learning-based combinatorial optimization approaches can be broadly grouped into three categories: constructive methods, improvement-based methods, and divide-and-conquer frameworks.

Constructive methods can be further divided into Local Construction (LC) and Global Prediction (GP) paradigms. LC approaches (Kool et al., 2019; Kwon et al., 2020; Kim et al., 2022; Berto et al., 2023; Drakulic et al., 2023; Pan et al., 2025) build solutions sequentially, selecting one variable at a time and ensuring feasibility at each step, but often suffer from myopic decisions and slow decoding. RL4CO community develops a comprehensive repository for this category of methods (Berto et al., 2023). GP approaches (Joshi et al., 2019; Fu et al., 2021; Qiu et al., 2022; Schuetz et al., 2022; Min et al., 2023; Sun & Yang, 2023; Xia et al., 2024; Li et al., 2024; Xiao et al., 2024), by contrast, predict global probability heatmaps in a single forward pass. This makes them efficient and globally aware, but the raw predictions often violate constraints and thus require post-processing to enforce feasibility. Within this paradigm, generative modeling methods (Hottung et al., 2021; Sun & Yang, 2023; Li et al., 2023; 2024; Zhao et al., 2024; Sanokowski et al., 2024) aim to learn a distribution over high-quality solutions for each instance. Feasible solutions can then be obtained by sampling from this distribution, which has been shown to yield competitive or even superior solver performance. To bridge these trade-offs, COExpander (Ma et al., 2025) recently introduced the *Adaptive Expansion* (AE) paradigm, which interpolates between LC’s fine-grained feasibility and GP’s global awareness. However, AE in COExpander is implemented only as a wrapper around GP predictors, leaving its effectiveness bounded by the quality and efficiency of the underlying backbone. This limitation motivates our work, which aims to instantiate AE as a native generative principle within a diffusion framework.

Improvement-based solvers (Chen & Tian, 2019; Oliveira da Costa et al., 2020; Wu et al., 2021; Sui et al., 2021; Li et al., 2021; Hou et al., 2023; Ma et al., 2023) focus on refining an initial solution, typically generated by a simple heuristic (e.g., greedy construction), through iterative optimization guided by neural networks or local search operators. While these methods can improve solution quality and naturally enforce feasibility at each step, they suffer from two key drawbacks: their dependence on heuristic initializers limits generality across problem settings, and the need for multiple refinement rounds incurs substantial computational overhead.

Divide-and-conquer (D&C) framewrks (Fu et al., 2021; Kim et al., 2021; Luo et al., 2023; 2025; Zheng et al., 2024; Ye et al., 2024) address scalability by decomposing large CO instances into smaller subproblems, solving them with either classical heuristics or neural solvers, and then aggregating the partial solutions into a global one. Scalability has long been a central challenge in neural CO: for GP-based methods, obtaining supervision signals is prohibitively expensive; for RL-based sequential models, issues like sparse rewards and unstable training further hinder their applicability to large-scale settings. DC provides a practical alternative to mitigate these limitations, and is largely orthogonal to constructive methods: it can be applied on top of LC, GP, or AE models, enabling them to operate effectively on larger-scale instances through problem decomposition.

A.2 MASK DIFFUSION MODEL

Diffusion models were first developed for continuous domains with Gaussian transitions (Sohl-Dickstein et al., 2015; Ho et al., 2020; Song et al., 2021), and later extended to discrete spaces by reformulating the forward process as a discrete-state Markov chain (Hoogeboom et al., 2021; Austin et al., 2021; Lou et al., 2024). Among these extensions, the *masked diffusion model* (MDM) has proven particularly effective: instead of arbitrary bit flips, variables are corrupted into an absorbing mask state, yielding semantically meaningful intermediate states. This simple yet powerful design enables principled training (Sahoo et al., 2024; Shi et al., 2024), scales efficiently to large models (Gong et al., 2025; Nie et al., 2025b; Wu et al., 2025; Sun et al., 2025; Ou et al., 2025), and consistently outperforms autoregressive models in reasoning and planning tasks (Ye et al., 2025; Nie et al., 2025a; Zheng et al., 2025). With these advantages, masked diffusion has become the de facto framework for discrete generative modeling, offering interpretable partial states, schedule-invariant

training, and efficient decoding that resonate strongly with the requirements of neural combinatorial optimization.

B ADDITIONAL EXPERIMENTS

B.1 CROSS-DISTRIBUTION AND CROSS-SCALE GENERALIZATION ON MISS

Table 7: Cross-distribution and cross-scale generalization results on MIS. Each entry reports the average solution size.

Algorithm	Cross-Distribution			Cross-Scale	
	$p = 0.2$	$p = 0.3$	$p = 0.4$	[350–400]	[1400–1600]
KaMIS	35.30	24.37	18.18	37.96	50.95
DIFUSCO	26.25	15.84	11.75	27.31	34.39
Fast T2T	29.52	17.77	13.27	32.56	36.95
NEXCO	31.24	18.14	14.17	34.65	38.39

Table 7 compares the performance of different solvers under distribution shifts and varying graph sizes. On random ER graphs with increasing edge probabilities ($p = 0.2, 0.3, 0.4$), learning-based solvers exhibit a clear gap compared to the heuristic KaMIS, but NEXCO consistently achieves larger solution sizes than both DIFUSCO and Fast T2T, e.g., 31.24 vs. 29.52 and 26.25 at $p = 0.2$. A similar trend holds in cross-scale settings: when transferring from RB-[350–400] and RB-[1400–1600], NEXCO again outperforms other learning-based solvers, achieving 34.65 and 38.39 respectively, compared to 32.56 and 36.95 from Fast T2T. These results highlight that native adaptive expansion maintains superior generalization under both distributional and scale shifts.

B.2 SCALABILITY ON TSP-10K

To further evaluate scalability, we conducted experiments on the large-scale TSP-10K dataset. We compare NEXCO with strong baselines, including the classical heuristic LKH, the divide-and-conquer approach GLOP (Ye et al., 2024), and diffusion-based solvers DIFUSCO (Sun & Yang, 2023), DISCO (Zhao et al., 2024), and Fast T2T (Li et al., 2024). For diffusion-based methods, we follow standard practice (Sun & Yang, 2023; Li et al., 2024) and apply K-Nearest Neighbor (KNN) sparsification with a sparse factor of 100 to restrict the search space by sampling 100 neighbors for each node. As shown in Table 8, NEXCO achieves the best balance between quality and efficiency: it reduces the optimality gap to 1.53% in 52s, improving over diffusion baselines (e.g., 1.63% in 70s by Fast T2T) while being substantially faster than LKH. These results highlight the scalability advantage of native adaptive expansion in large-scale settings.

To extend to larger regimes, NEXCO is naturally compatible with standard scaling strategies used in large-graph CO systems, such as divide-and-conquer frameworks (e.g., region partitioning similar to GLOP (Ye et al., 2024)), or replacing the backbone GNN with lightweight sparse-attention Transformers (Luo et al., 2025). These directions do not alter the proposed generative principle and can be incorporated in future work to achieve industrial-scale deployments. In addition, the partial-solution semantics of NEXCO provide a natural mechanism for overcoming the scarcity and high cost of supervision in large-scale CO tasks. Because the model operates on feasible partial states rather than requiring complete high-quality solutions at every step, it can be trained using incomplete, heuristic, or low-cost labels and gradually refine its own predictions through iterative self-training. This significantly reduces reliance on exact solvers, whose computational cost grows prohibitively with instance size. The ability to learn from weak or partial supervision aligns with recent scalable training pipelines (Luo et al., 2025; Li et al., 2025) that couple lightweight backbones with progressive bootstrapping. Incorporating such strategies offers a promising path for applying NEXCO in settings where high-quality ground-truth labels are expensive or unavailable, thereby further extending its applicability to truly large and industrial-scale CO environments.

864
 865 Table 8: **Results on TSP-10K.** NEXCO achieves the best trade-off between solution quality and
 866 runtime.

Method	Type	Length ↓	Drop ↓	Time
LKH	Heuristics	71.77	–	332s
GLOP (more revisions)	DC	75.29	4.90%	15s
DIFUSCO ($T_s = 100$)	GP	73.91	2.98%	124s
DISCO	GP	73.84	2.88%	92s
Fast T2T ($T_s = 5$)	GP	72.94	1.63%	70s
NEXCO ($D_s = 5$)	NAE	72.87	1.53%	52s

874
 875 Table 9: Solution quality (%) for methods trained on TSP-100 problems and evaluated on **TSPLIB**
 876 instances with 50–200 nodes. * denotes results quoted from previous works (Li et al., 2024).

Instances	AM*	GCN*	Learn2OPT*	GNNGLS*	DIFUSCO*	T2T*	Fast T2T*	NEXCO
eil51	16.767	40.025	1.725	1.529	2.82	0.14	0.00	0.00
berlin52	4.169	33.225	0.449	0.142	0.00	0.00	0.00	0.00
st70	1.737	24.785	0.040	0.764	0.00	0.00	0.00	0.00
eil76	1.992	27.411	0.096	0.163	0.34	0.00	0.00	0.00
pr76	0.816	27.793	1.228	0.039	1.12	0.40	0.00	0.00
rat99	2.645	17.633	0.123	0.550	0.09	0.09	0.00	0.00
kroA100	4.017	28.828	18.313	0.728	0.10	0.00	0.00	0.00
kroB100	5.142	34.686	1.119	0.147	2.29	0.74	0.65	0.00
kroc100	0.972	35.506	0.349	1.571	0.00	0.00	0.00	0.00
krod100	2.717	38.018	0.866	0.572	0.07	0.00	0.00	0.00
kroE100	1.470	26.589	1.832	1.216	3.83	0.27	0.00	0.00
rd100	3.407	30.432	1.725	0.003	0.08	0.00	0.00	0.00
eil101	2.994	26.701	0.387	1.529	0.03	0.00	0.00	0.00
lin105	1.739	34.902	1.867	0.606	0.00	0.00	0.00	0.54
pr107	3.933	30.564	0.898	0.439	0.91	0.61	0.62	0.08
pr124	3.677	70.146	10.322	0.755	1.02	0.60	0.08	0.58
bier127	5.908	45.561	3.044	1.948	0.94	0.54	1.50	0.66
ch130	3.182	39.090	0.709	3.519	0.29	0.06	0.00	0.00
pr136	5.064	58.673	0.000	3.387	0.19	0.10	0.01	0.00
pr144	7.641	55.837	1.526	3.581	0.80	0.50	0.39	0.00
ch150	4.584	49.743	0.312	2.113	0.57	0.49	0.00	0.00
kroA150	3.784	45.411	0.724	2.984	0.34	0.14	0.00	0.00
kroB150	2.437	36.745	0.886	3.258	0.30	0.00	0.07	0.02
pr152	7.494	33.925	0.029	3.119	1.69	0.83	1.19	0.68
u159	7.551	38.338	0.054	1.020	0.82	0.00	0.00	0.00
rat195	6.839	24.968	0.743	1.666	1.48	1.27	0.79	0.11
d198	373.020	62.351	0.522	4.772	3.32	1.97	0.86	0.00
kroA200	7.106	40.885	1.441	2.029	2.28	0.57	0.49	0.00
kroB200	8.541	43.643	2.064	2.589	2.35	0.92	2.50	0.00
Mean	16.767	40.025	1.725	1.529	0.97	0.35	0.28	0.09

903
 904
 905 **B.3 EVALUATION ON REAL-WORLD INSTANCES**
 906

907 We evaluate our model on real-world TSPLIB instances with 50–200 nodes (Reinelt, 1991). The
 908 model is trained on TSP-100 dataset, and compared against state-of-the-art baselines including DI-
 909 FUSCO (Sun & Yang, 2023), T2T (Li et al., 2023), and Fast T2T (Li et al., 2024). The hyperpa-
 910 rameter configurations for the diffusion-based baselines are as follows: DIFUSCO with $T_s = 50$;
 911 T2T with $T_s = 50$ and $T_g = 30$; Fast T2T (with guided sampling) with $T_s = 10$ and $T_g = 10$, and
 912 NEXCO with $D_s = 20$. All diffusion-based methods are evaluated under the same settings, using
 913 greedy decoding followed by 2-Opt local search as post-processing. For consistency, the coordinates
 914 of each TSPLIB instance are normalized to the range $[0, 1]$.

915
 916
 917 **B.4 ABLATION STUDY ON THE CORRUPTION SCHEME**

We compare our mask-based forward corruption with a uniform corruption scheme that treats all variables as equally likely to be flipped, similar to the uniform perturbation used in Fast-T2T (Table 10). Across TSP-500 and TSP-1000, and under the same number of expansion steps, the mask-based model consistently achieves lower optimality gaps. The difference arises from the structural behavior of the two corruption processes. Uniform corruption injects noise indiscriminately and often produces partial states with weak or noisy supervision signals for the denoiser. In contrast, the proposed mask corruption applies one-way absorbing $1 \rightarrow 0$ updates that preserve the feasible structure of partial solutions while selectively revealing informative variables. This yields clearer denoising targets and a more stable reverse trajectory, which in turn explains the consistent improvements observed in our ablations.

Table 10: Performance comparison across different corruption scheme.

Exp. Step	Scheme	TSP-500			TSP-1000		
		Obj.	Gap \downarrow	Time	Obj.	Gap \downarrow	Time
3	Uniform	16.66	0.65%	0.27s	23.34	0.94%	1.00s
	Mask	16.61	0.39%	0.23s	23.31	0.85%	0.91s
5	Uniform	16.65	0.59%	0.39s	23.32	0.86%	1.44s
	Mask	16.59	0.28%	0.33s	23.26	0.63%	1.31s
7	Uniform	16.63	0.53%	0.49s	23.30	0.78%	1.85s
	Mask	16.59	0.25%	0.43s	23.24	0.52%	1.68s

B.5 ABLATION STUDY ON THE ADAPTIVE EXPANSION

Table 11 reports the comparison between NAE and a non-adaptive “global t-schedule” baseline. The global baseline does not perform any form of adaptive expansion. It runs the diffusion model for a fixed number of denoising steps, produces a dense full prediction at the final step, and then applies a single greedy decoding to obtain a complete solution. No intermediate partial-state construction is carried out, and the amount of expansion is fixed rather than guided by model confidence. In contrast, NAE operates directly on feasible partial states and expands them progressively. At each step, it activates candidates according to their confidence scores and applies feasibility projection to maintain monotone growth of the partial solution. This native adaptivity enables the model to commit early to high-confidence regions while deferring uncertain components to later steps, thereby structuring the decoding trajectory in a principled way. As shown in the ablation, this constructive and confidence-aware expansion yields significantly smaller optimality gaps than the global non-adaptive schedule under the same number of denoiser calls. The improvement therefore stems not from additional computation, but from the design of an adaptive partial-state expansion mechanism.

Table 11: Comparison between Global and NAE under different expansion steps.

Exp. Step	Method	TSP-500			TSP-1000		
		Obj.	Gap \downarrow	Time (s)	Obj.	Gap \downarrow	Time (s)
3	Global	16.67	0.73%	0.22s	23.39	1.16%	0.81s
	NAE	16.61	0.39%	0.23s	23.31	0.85%	0.91s
5	Global	16.67	0.71%	0.29s	23.34	0.96%	1.04
	NAE	16.59	0.28%	0.33s	23.26	0.63%	1.31s
7	Global	16.65	0.63%	0.36s	23.33	0.92%	1.29s
	NAE	16.59	0.25%	0.43s	23.24	0.52%	1.68s

B.6 ABLATION STUDY ON THE FEASIBILITY PROJECTION

972
 973 Table 12 compares our feasibility projection mechanism with a greedy decoding baseline. In the
 974 greedy baseline, a complete solution is constructed at every diffusion step by selecting variables
 975 according to their predicted probabilities. After obtaining this full solution, a fixed proportion of
 976 variables is remasked, and the resulting state is used as input to the next diffusion step. This produces
 977 a wrapper-like refinement cycle that repeatedly rebuilds full solutions throughout the trajectory.

978 In contrast, our projection operator Γ maintains a single monotone partial-solution trajectory. At
 979 each step, Γ accepts a candidate activation only when feasibility is preserved, and it does not generate
 980 full solutions prematurely. This prevents the repeated reconstruction inherent to greedy decoding
 981 and avoids the error accumulation introduced by successive remasking cycles. Empirically, both ap-
 982 proaches can eventually achieve similar optimality gaps when sufficient steps are allowed. However,
 983 greedy decoding consistently incurs higher runtime and exhibits less stable behavior due to its re-
 984 construct–remask procedure. These results demonstrate that the feasibility projection used by NAE
 985 offers a more efficient and principled alternative to wrapper-style greedy refinement strategies.

986
 987 Table 12: Comparison between Greedy decoding and Projection under different expansion steps.

988 989 Exp. Step	990 Method	991 TSP-500			992 TSP-1000		
		993 Obj.	994 Gap \downarrow	995 Time (s)	996 Obj.	997 Gap \downarrow	998 Time (s)
999 1000 3	Greedy decoding	16.61	0.39%	0.39s	23.33	0.94%	1.29s
	Projection	16.61	0.39%	0.23s	23.31	0.85%	0.91s
1001 5	Greedy decoding	16.60	0.31%	0.54s	23.26	0.63%	1.95s
	Projection	16.59	0.28%	0.33s	23.26	0.63%	1.31s
1002 7	Greedy decoding	16.59	0.26%	0.71s	23.25	0.57%	2.59s
	Projection	16.59	0.25%	0.43s	23.24	0.52%	1.68s

1003
 1004 **B.7 RUNTIME BREAKDOWN OF NAE INFERENCE**

1005
 1006 In this section, we provide detailed runtime profiling of NAE across all tasks and instance sizes.
 1007 Table 13 reports total inference time per instance, together with the proportion attributed to denoiser
 1008 evaluations and to the projection operator $\Gamma(\cdot)$.

1009
 1010 Table 13: Runtime breakdown of NAE inference across all tasks.

1011 Task	1012 Total Time (s)	1013 Denoiser	1014 $\Gamma(\cdot)$ Projection
1015 TSP100	0.018	0.016 (88.89%)	0.002 (11.11%)
1016 TSP500	0.097	0.081 (83.51%)	0.016 (16.49%)
1017 TSP1000	0.218	0.162 (74.31%)	0.056 (25.69%)
1018 MIS–RBsmall	0.033	0.026 (78.79%)	0.007 (21.21%)
1019 MIS–RBlarge	0.190	0.173 (91.05%)	0.017 (8.95%)
1020 MIS–ER	0.200	0.184 (92.00%)	0.016 (8.00%)
1021 CVRP50	0.022	0.017 (77.27%)	0.005 (22.73%)
1022 CVRP100	0.030	0.021 (70.00%)	0.009 (30.00%)
1023 CVRP200	0.078	0.052 (66.67%)	0.026 (33.33%)

1024
 1025 Across all settings, denoiser calls constitute the primary computational bottleneck, consistently ac-
 1026 counting for the majority of inference time. The relative cost of the projection operator $\Gamma(\cdot)$ increases
 1027 with instance size due to the larger number of candidate variables that must be feasibility-checked.
 1028 However, its absolute cost remains small (on the order of milliseconds), and the overall inference
 1029 complexity of NAE continues to be dominated by denoiser evaluations. These results demonstrate
 1030 that the projection mechanism introduces only a lightweight overhead that scales predictably with
 1031 problem size, confirming that NAE’s efficiency is primarily driven by its denoiser architecture rather
 1032 than by feasibility projection.

1026 **C MODEL ARCHITECTURE DETAILS**
 1027

1028 **C.1 INPUT EMBEDDING LAYER**
 1029

1030 Given node vector $x \in \mathbb{R}^{N \times 2}$, weighted edge vector $e \in \mathbb{R}^E$, denoising timestep $t \in \{\tau_1, \dots, \tau_M\}$,
 1031 where N denotes the number of nodes in the graph, and E denotes the number of edges, we compute
 1032 the sinusoidal features of each input element respectively:

$$\tilde{x}_i = \text{concat}(\tilde{x}_{i,0}, \tilde{x}_{i,1}), \quad (9)$$

$$\tilde{x}_{i,j} = \text{concat}\left(\sin \frac{x_{i,j}}{T^{0/d}}, \cos \frac{x_{i,j}}{T^{0/d}}, \sin \frac{x_{i,j}}{T^{2/d}}, \cos \frac{x_{i,j}}{T^{2/d}}, \dots, \sin \frac{x_{i,j}}{T^{d/d}}, \cos \frac{x_{i,j}}{T^{d/d}}\right), \quad (10)$$

$$\tilde{e}_i = \text{concat}\left(\sin \frac{e_i}{T^{0/d}}, \cos \frac{e_i}{T^{0/d}}, \sin \frac{e_i}{T^{2/d}}, \cos \frac{e_i}{T^{2/d}}, \dots, \sin \frac{e_i}{T^{d/d}}, \cos \frac{e_i}{T^{d/d}}\right), \quad (11)$$

1033 where d is the embedding dimension, T is a large number (usually selected as 10000), and $\text{concat}(\cdot)$
 1034 denotes concatenation. In CVRP, each node is described not only by its coordinates but also by
 1035 customer demand c_i and a depot indicator $\delta_i \in \{0, 1\}$. We embed these heterogeneous features
 1036 separately and merge them into the node representation:

$$\tilde{c}_i = \text{concat}\left(\sin \frac{c_i}{T^{0/d}}, \cos \frac{c_i}{T^{0/d}}, \dots, \sin \frac{c_i}{T^{d/d}}, \cos \frac{c_i}{T^{d/d}}\right), \quad (12)$$

$$\tilde{\delta}_i = \text{Embed}_{\text{depot}}(\delta_i), \quad (13)$$

$$\tilde{x}_i = \tilde{x}_i + \tilde{d}_i + \tilde{\delta}_i. \quad (14)$$

1043 Next, we compute the input features of the graph convolution layer:
 1044

$$x_i^0 = W_1^0 \tilde{x}_i, \quad (15)$$

$$e_i^0 = W_2^0 \tilde{e}_i. \quad (16)$$

1045 Specifically, for TSP and CVRP, the embedding input edge vector e is a weighted adjacency matrix,
 1046 which represents the distance between different nodes, and e^0 is computed as above. For MIS, we
 1047 initialize e^0 to a zero matrix $0^{E \times d}$.

1048 **C.2 GRAPH CONVOLUTION LAYER**
 1049

1050 Following (Joshi et al., 2019), the cross-layer convolution operation is formulated as:
 1051

$$x_i^{l+1} = x_i^l + \text{ReLU}(\text{BN}(W_1^l x_i^l + \sum_{j \sim i} \eta_{ij}^l \odot W_2^l x_j^l)), \quad (17)$$

$$e_{ij}^{l+1} = e_{ij}^l + \text{ReLU}(\text{BN}(W_3^l e_{ij}^l + W_4^l x_i^l + W_5^l x_j^l)), \quad (18)$$

$$\eta_{ij}^l = \frac{\sigma(e_{ij}^l)}{\sum_{j' \sim i} \sigma(e_{ij'}^l) + \epsilon}, \quad (19)$$

1052 where x_i^l and e_{ij}^l denote the node feature vector and edge feature vector at layer l , $W_1, \dots, W_5 \in$
 1053 $\mathbb{R}^{h \times h}$ denote the model weights, and η_{ij}^l denotes the dense attention map. The convolution operation
 1054 integrates the edge feature to accommodate the significance of edges in routing problems.

1055 For TSP and CVRP, we aggregate the edge convolutional feature and reformulate the update for
 1056 edge features as follows:

$$e_{ij}^{l+1} = e_{ij}^l + \text{ReLU}(\text{BN}(W_3^l e_{ij}^l + W_4^l x_i^l + W_5^l x_j^l)) + W_6^l(\text{ReLU}(t^0)). \quad (20)$$

1057 For MIS, we aggregate the node convolutional feature and reformulate the update for node features
 1058 as follows:

$$x_i^{l+1} = x_i^l + \text{ReLU}(\text{BN}(W_1^l x_i^l + \sum_{j \sim i} \eta_{ij}^l \odot W_2^l x_j^l)) + W_6^l(\text{ReLU}(t^0)). \quad (21)$$

1080
1081

C.3 OUTPUT LAYER

1082

The prediction of the edge heatmap in TSP and CVRP, and node heatmap in MIS is as follows:

1083

$$e_{i,j} = \text{Softmax}(\text{norm}(\text{ReLU}(W_e e_{i,j}^L))), \quad (22)$$

1084

$$x_i = \text{Softmax}(\text{norm}(\text{ReLU}(W_n x_i^L))), \quad (23)$$

1085

where L is the number of GCN layers and norm is layer normalization.

1086

1087

C.4 MODEL PARAMETERS

1088

1089

For all tasks, we adopt a 12-layer GCN as described above. For TSP, following the setting of (Sun & Yang, 2023), we apply a K-Nearest Neighbor (KNN) strategy to sparsify the graph in order to reduce training memory and shrink the search space. Specifically, for TSP-500 and TSP-1000, the sparsity factors are set to 50 and 100, respectively.

1090

1091

D EXPERIMENTAL SETUP

1092

1093

D.1 HARDWARE

1094

1095

1096

1097

1098

All models are trained and tested using NVIDIA A40 (48G) GPUs and Intel(R) Xeon(R) Gold 5220 CPU @ 2.20GHz. All test evaluations are performed in a single-threaded setting, where the average runtime per instance is reported to ensure fair comparison across different models.

1099

1100

We have organized the training settings and model parameters of NEXCO in Table 14. For all problems, we adopt a curriculum learning strategy, where models are progressively fine-tuned from smaller datasets to large ones.

1101

1102

D.2 TRAINING SETUP

1103

1104

1105

1106

1107

1108

1109

Table 14: Details about the training hyperparameters of NEXCO.

Problem	Data	Data Size	Batch Size	Epoch	Learning Rate	Hidden Dimension
TSP	Uniform-100	1,502k	16	100	2e-4	256
TSP	Uniform-500	128k	6	50	2e-4	256
TSP	Uniform-1000	64k	4	50	2e-4	256
MIS	RB-200-300	90k	4	50	2e-4	256
MIS	RB-800-1200	6.4k	1	10	5e-5	256
MIS	ER-700-800	163k	4	50	2e-4	128
CVRP	Uniform-50	1,280k	32	50	2e-4	256
CVRP	Uniform-100	640k	12	50	2e-4	256
CVRP	Uniform-200	32k	2	50	2e-4	256

1120

1121

1122

E LICENSES

1123

The licenses for the codes used in this work are listed in Table 15.

1124

1125

1126

1127

F LLM USAGE

1128

1129

1130

1131

1132

1133

1134

1135

1136

1137

1138

1139

1140

1141

1142

1143

1144

1145

1146

1147

1148

1149

1150

1151

1152

1153

1154

1155

1156

1157

Table 15: Licenses for codes used in this work

Resource	Type	Link	License
LKH3 (Helsgaun, 2017)	Code	http://webhotel4.ruc.dk/~keld/research/LKH-3/	Available for academic research
HGS (Vidal et al., 2012)	Code	https://github.com/chkwon/PyHgese	MIT License
Concorde (Applegate et al., 2006)	Code	https://github.com/jvkersch/pyconcorde	BSD 3-Clause License
KaMIS (Lamm et al., 2016)	Code	https://github.com/KarlsruherMIS/KaMIS	MIT License
AM (Kool et al., 2019)	Code	https://github.com/wouterkool/attention-learn-to-route	MIT License
BQ-NCO (Drakulic et al., 2023)	Code	https://github.com/naver/bq-nco	CC BY-NC-SA 4.0
GCN (Joshi et al., 2019)	Code	https://github.com/chaitjo/graph-convnet-tsp	MIT License
DIMES (Qiu et al., 2022)	Code	https://github.com/DIMESteam/DIMES	MIT License
DIFUSCO (Sun & Yang, 2023)	Code	https://github.com/Edward-Sun/DIFUSCO	MIT License
T2T (Li et al., 2023)	Code	https://github.com/Thinklab-SJTU/T2TCO	MIT License
Fast T2T (Li et al., 2024)	Code	https://github.com/Thinklab-SJTU/Fast-T2T	MIT License
COExpander (Ma et al., 2025)	Code	https://github.com/Thinklab-SJTU/COExpander	MIT License
LEHD (Luo et al., 2023)	Code	https://github.com/CIAM-Group/NCO_code/tree/main/single_objective/LEHD	Not specified
Sym-NCO (Kim et al., 2022)	Code	https://github.com/alstn12088/Sym-NCO	MIT License
			Not specified

1166

1167

1168

1169

1170

1171

1172

1173

1174

1175

1176

1177

1178

1179

1180

1181

1182

1183

1184

1185

1186

1187