

Deep Learning for Optimization of Trajectories for Quadrotors

Yuwei Wu, Xiaotao Sun, Igor Spasojevic, Vijay Kumar

Abstract—This paper presents a novel learning-based trajectory planning framework for quadrotors that combines model-based optimization techniques with deep learning. Specifically, we formulate the trajectory optimization problem as a quadratic programming (QP) problem with dynamic and collision-free constraints using piecewise trajectory segments through safe flight corridors [1]. We train neural networks to directly learn the time allocation for each segment to generate optimal smooth and fast trajectories. Furthermore, the constrained optimization problem is applied as a separate implicit layer for backpropagation in the network, for which the differential loss function can be obtained. We introduce an additional penalty function to penalize time allocations which result in solutions that violate the constraints to accelerate the training process and increase the success rate of the original optimization problem. To this end, we enable a flexible number of sequences of piece-wise trajectories by adding an extra end-of-sentence token during training. We illustrate the performance of the proposed method via extensive simulation and experimentation and show that it works in real time in diverse, cluttered environments.

I. INTRODUCTION

The application of micro aerial vehicles (MAVs) has significantly expanded recently in agriculture, industrial inspection, and customer services [2]. As one crucial module of autonomy, the research of trajectory generation has made significant progress in realistic feasibility and computation efficiency. Traditional optimization-based trajectory generation with simplified dynamics offers computational advantages for real-time deployments [3] but suffers from parameter tunings and infeasible formulations when applied to real-world data. Moreover, existing learning-based methods [4]–[6] are mainly trained to generate commands or waypoints, limited by dynamic interpretability and fail to provide guarantees for real-time safe flight. This paper addresses online trajectory generation, focusing on achieving efficiency, reliability, and optimality.

There is extensive work on optimization-based approaches that rely on approximating the free space by safe-flight corridors which are defined by a set of overlapping convex subsets. These approaches find the optimal trajectory by solving a quadratic program that assumes a specified *time allocation*, an assignment of time intervals to each convex subset [3]. However, there are no principled approaches to solving the time allocation problem. Most previous work decouples the time allocation problem from the optimization problem and focuses on refining and scaling the time with heuristics to

The authors are with the GRASP Laboratory, University of Pennsylvania, Philadelphia, PA, 19104 USA {yuweiwu, sxt, igorspas, kumar}@seas.upenn.edu.

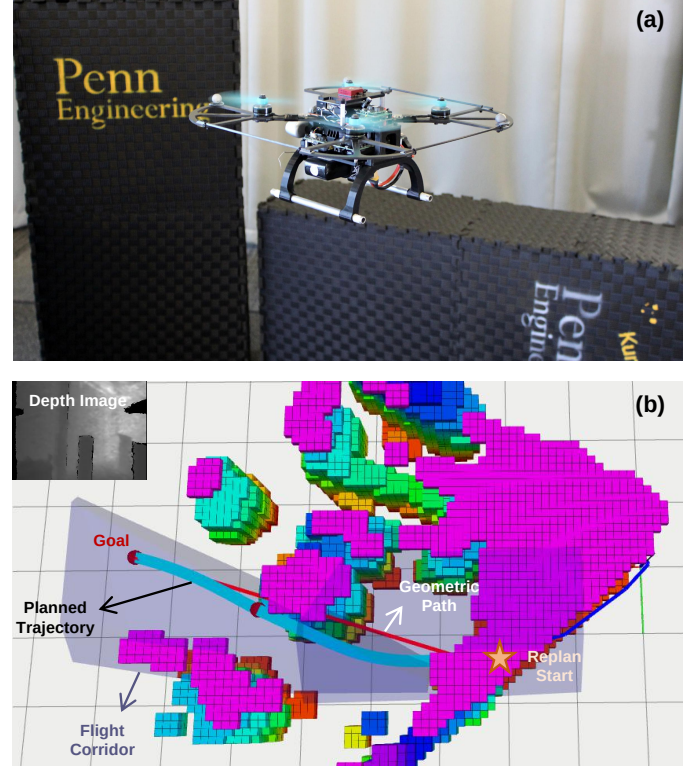


Fig. 1. (a) Real-time deployment on Falcon 250 v2 autonomous stack. (b) The planned trajectory (in light blue) demonstration with time allocation. We use a safe flight corridor (in purple) and a geometrical searched path (in red) and allocate time by a neural network. Then the trajectory generation is solved by fixed-time quadratic programming. The video is available at: <https://youtu.be/tA02dJz9ux8>

arrive at a feasible time allocation [7]–[15]. Parameterizing the optimal time is computationally expensive, and even so, does not provide optimality guarantees. Including time parametrization results in a joint optimization of piecewise trajectories and time intervals, leading to suboptimality. Most state-of-the-art solvers cannot find feasible solutions with a reasonable computational budget. These works inspired us to explore the nature of trajectory generation as an optimization problem and improve the efficiency and optimality of the existing methods.

Leveraging the advantages of both model-based and learning-based methods, we propose an efficient learning-based planning framework that utilizes neural networks to solve the optimal time allocation problem for each minimum control trajectory. The trajectory generation problem is formulated by finding the optimal trajectory coefficients with discrete dynamic constraints and loose constraints for intermediate waypoints in overlapping regions of flight corri-

dors. We present a neural network to output time allocation, with an extra differential implicit layer [16] solving trajectory optimization with hard constraints parameterized by time allocations. In addition, we design loss functions related to the original nonlinear programming while considering different solving results and adding infeasible penalty terms to train the network efficiently. We summarize our contributions as follows:

- A neural network architecture incorporates trajectory optimization problems as implicit layers with a flexible sequence of trajectory segments.
- An interpretable and lightweight learning-based planning framework that directly learns optimal time allocation from the corridor and state constraints for minimum control trajectory generations with quadratic programming.
- Real-world deployment of our proposed method on hardware platform integrating on-board autonomy and high-performance demonstration in both indoor and outdoor environments, shown in Fig. 1. The open source code is available at: <https://github.com/KumarRobotics/AllocNet>.

II. RELATED WORKS

A. Optimal Time Allocation for Trajectory Planning

Decoupled trajectory parameterization of flat outputs with polynomials has been well-explored as an efficient way to optimize the trajectory with quadratic programming (QP) while encoding dynamic and obstacle-free constraints with differential flat outputs [3, 12]. A closed-form solution for this QP formulation, incorporating equality constraints, was proposed in [12]. Its numerical stability issues during matrix inverse were addressed in [17], and better-conditioned matrices are formulated in [18] with linear computational complexity to solve QP. However, introducing time allocation optimization for piece-wise trajectory makes the problem more challenging, resulting in a joint spatial-temporal ill-conditioned nonlinear programming (NLP). Previous research has primarily focused on iterative time refinement [19], time-optimal primitives [20, 21], path parameterization [13, 22], or bi-level optimization [14] to address the challenge by solving the problem in two phases or iteratively searching for optimal time allocation. Moreover, different analytic gradients towards time and waypoints have been studied [11] to improve the efficiency of fast trajectory generation. However, the lack of feasibility guarantees and the high sensitivity of parameters limit the performance of these approaches in real-world scenarios.

B. Learning for Autonomous Navigation

Learning-based methods have emerged as powerful approaches for autonomous navigation applications. End-to-end policy learning [4] directly trained the control commands using deep reinforcement learning to generate trajectory with waypoints. Many existing works assume a collision-free reference line and perform policy learning based on state observations along this predefined path or in an obstacle-free space. These methods involve imitating human drivers or expert planners [5, 6] or utilizing learning-based perception to generate collision-free waypoints [23]. However, they often

generate over-simplified end-to-end commands or waypoints that require additional components such as Model Predictive Control (MPC) or a dynamic trajectory planner to obtain a smoother trajectory [24]–[26]. Differentiable trajectory optimization mapping from waypoints can also be included during the learning process [26]. However, the computational complexity associated with the learning processes and generalization limits the deployment to more unknown scenarios. Addressing time allocation in trajectory optimization, Ryou et al. [27] employed a deep neural network with sequence-to-sequence learning to determine optimal time allocation given waypoints. Nevertheless, the complex process of multi-phase learning with brute-search time allocations remains a challenge that must be addressed, encouraging us to pursue a lightweight and efficient structure.

C. Neural Optimizations with Hard Constraints

Recent works using neural networks (NNs) to learn optimizers with hard constraints have gained significant attention due to their extensive applications and promising results. Amos et al. [16, 28] proposed OptNet to demonstrate the effective integration of quadratic programming (QP) with hard constraints within the NNs framework, which illustrate the potential of leveraging implicit layers to tackle complex optimization problems. Jaquier et al. [29] extended the application of OptNet to multitask control problems by formulating these tasks as QPs and employing differentiable layers to enable the learning of specific manipulator skills. However, these approaches primarily focus on feasible problem formulations, leaving a gap when dealing with the inherent infeasibility of specific optimization problems. To address this limitation, DC3 [30] fixed the violation and obtained an approximated solution by adding additional differentiable completion and correction procedures, which can be extended to nonlinear optimization. Other works using meta-learning [31], double descriptions [32] can also be applied to learn optimizers. Integrating differential optimization in NNs provides diverse tools to address optimization problems.

III. DEEP TIME-OPTIMAL TRAJECTORY LEARNING

A. Spatial-temporal Trajectory Optimization

Leveraging the differential flatness of quadrotor dynamics, we represent the trajectory of the quadrotor by the trajectory of its flat outputs: $\sigma = [x, y, z, \psi]^T \in \mathbb{R}^3 \times S^1$. The vector $[x, y, z]^T \in \mathbb{R}^3$ represents the position of the quadrotor, whereas $\psi \in S^1$ represents the “yaw” angle that the projection of its body x axis forms with the x axis of the world frame. The remarkable fact is that a generic trajectory of the quadrotor corresponds to a unique function σ , and in turn, every sufficiently smooth σ corresponds to a dynamically feasible trajectory of the vehicle, *provided its motors can exert sufficient thrust*. The latter, actuation constraints render numerous classes of trajectory optimization problems for quadrotors computationally challenging. In this work, we focus on the class of minimum-time problems, enforcing the constraints in an approximate discrete approach on the derivatives of σ , in the hope of developing an algorithm that

can output near-optimal and *safe* (collision-free, dynamically feasible) trajectories in real-time. We parameterize the trajec-

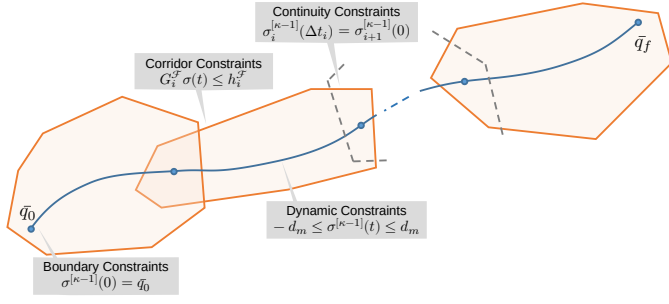


Fig. 2. Trajectory Optimization formulation with corridors and dynamic constraints. The equality boundary and continuity constraints are also demonstrated.

tory of flat outputs $\sigma(t)$ as a suitably smooth $N = 2\kappa - 1$ degree piece-wise polynomial function. In particular, $\sigma(t)$ is a concatenation of a sequence of $M \in \mathbb{N}$ segments, $\sigma_1(\cdot), \sigma_2(\cdot), \dots, \sigma_M(\cdot)$. Segment i , σ_i , is characterized by its coefficients $\mathbf{c}_i \in \mathbb{R}^{(N+1) \times 3}$ and duration $\Delta t_i > 0$. The complete trajectory $\sigma(t)$ is therefore encoded through its coefficient matrix $\mathbf{c} = [\mathbf{c}_1^T, \dots, \mathbf{c}_M^T]^T$ and time allocation intervals $\mathbf{t} = [\Delta t_1, \dots, \Delta t_M]^T$, with:

$$\sigma(t) = \sigma_i(t - \sum_{k=1}^{i-1} \Delta t_k), \quad \forall t \in [\sum_{k=1}^{i-1} \Delta t_k, \sum_{k=1}^i \Delta t_k], \quad (1)$$

where $\sigma_i(t) = \mathbf{c}_i^T \beta(t)$,
 $\beta(t) = [1, t, t^2, \dots, t^N]^T, \quad \forall t \in [0, \Delta t_i].$

Given a sequence of safe flight corridors extracted from free space, together with dynamic feasibility in the form of collision-free constraints on the derivatives of the flat outputs, generate the time-optimal vector of duration intervals and corresponding coefficients. For example, “collision-free” constraints on the second derivatives encode actuation bounds on the acceleration of the vehicle, constraints on the first derivatives encode bounds on its velocity, and constraints on the zeroth derivatives encode staying in the obstacle-free region of space, shown in Fig. 2.

The objective function involves minimizing the traversal time as well as the control effort (snap or jerk) required to execute the planned n -dimension piece-wise polynomial trajectory. Hence, the problem is formulated by

$$\min_{\mathbf{c}, \mathbf{t}} \int_0^{\sum_{k=1}^M \Delta t_k} \|\sigma^{(\kappa)}(t)\|^2 dt + w_t f(\mathbf{t}), \quad (2a)$$

$$\text{s.t.} \quad \sigma^{[\kappa-1]}(0) = \bar{\mathbf{q}}_0, \quad \sigma^{[\kappa-1]}(\sum_{k=1}^M \Delta t_k) = \bar{\mathbf{q}}_f, \quad (2b)$$

$$\sigma_i^{[\kappa-1]}(\Delta t_i) = \sigma_i^{[\kappa-1]}(0), \quad (2c)$$

$$-d_m \leq \sigma_i(t)^{[\kappa-1]} \leq d_m, \quad (2d)$$

$$\mathbf{G}_i^H \sigma_i(t) \leq \mathbf{h}_i^H, \quad (2e)$$

$$\forall i \in 1, \dots, M, \quad \forall t \in [0, \Delta t_i]. \quad (2f)$$

The $\bar{\mathbf{q}}_0, \bar{\mathbf{q}}_f \in \mathbb{R}^{3 \times \kappa}$ are initial and end states up to $(\kappa - 1)$ -order derivative and $\sigma^{[\kappa-1]}(t) = [\sigma(t)^T, \dots, \sigma^{(\kappa-1)}(t)^T]^T$.

The time objective $f(t)$ could be the sum of total traversal time or its absolute deviation from an expected finish time, and w_t is the associated weight. We use Eq. (2d) to include all state and input inequality constraints except the constraints on position, which are captured in Eq. (2e). The vector $\mathbf{d}_m \in \mathbb{R}^{\kappa-1}$ includes the magnitude of maximum velocity v_{max} and acceleration a_{max} for $k = 3$, adding maximum jerk j_{max} for $k = 4$. The corridor constraints are formulated by half-space representation of polytopes with $\mathbf{G}_i^H \in \mathbb{R}^{F_i \times 3}$, $\mathbf{h}_i^H \in \mathbb{R}^{F_i}$, where F_i is the number of hyperplanes for i -th polytope. The continuous inequality constraints are enforced by discretizing the trajectory by the sampling number N_{res} and querying sample points to satisfy positions within corridors and higher derivatives within maximum limits. More generally, we can formulate the problem into nonlinear programming as follows.

Problem 1. *The Time-optimal minimum control trajectory optimization with corridor and dynamics constraints is a parametric nonlinear programming (NLP) problem.*

$$\min_{\mathbf{c}, \mathbf{t}} \mathbf{c}^T \mathbf{Q}(\mathbf{t}) \mathbf{c} + w_t f(\mathbf{t}), \quad (3a)$$

$$\text{s.t.} \quad \mathbf{A}(\mathbf{t}) \mathbf{c} = \mathbf{b}, \quad (3b)$$

$$\mathbf{G}(\mathbf{t}) \mathbf{c} \leq \mathbf{h}, \quad (3c)$$

$$\mathbf{H} \mathbf{t} = \mathbf{m}, \quad (3d)$$

$$\mathbf{C} \mathbf{t} \leq \mathbf{d}, \quad (3e)$$

where $\mathbf{Q} \succeq \mathbf{0}$ is a positive semidefinite matrix.

The former problem presents a non-convex optimization that nevertheless possesses a certain amount of structure. In particular, given a feasible time allocation, we can directly compute the time-dependent matrices $\mathbf{Q}_t, \mathbf{A}_t, \mathbf{G}_t$ and solve for the optimal vector of coefficients as a convex *quadratic* programming. Importantly, all aforementioned constraints are *linear* constraints on \mathbf{c} for a fixed \mathbf{t} .

Problem 2. *The minimum control trajectory optimization for a given time allocation can be formulated as a quadratic programming (QP) problem as follows:*

$$\min_{\mathbf{c}} \mathbf{c}^T \mathbf{Q}_t \mathbf{c}, \quad (4a)$$

$$\text{s.t.} \quad \mathbf{A}_t \mathbf{c} = \mathbf{b}, \quad (4b)$$

$$\mathbf{G}_t \mathbf{c} \leq \mathbf{h}. \quad (4c)$$

Unfortunately, the *quadratic* problem is not always feasible with corridors and dynamics constraints for a specific time allocation, which introduces extra refinements and iterations for better time intervals.

B. Learning Time Allocation with NNs

1) *Network architecture:* We propose a Time Allocation Network (*AllocNet*) to learn the time allocation for trajectory generation. The full network architecture and system framework are shown in Fig. 3. We learn a function that maps the flattened initial and terminal states, denoted by $\mathbf{q} = [\bar{\mathbf{q}}_0, \bar{\mathbf{q}}_f]^T$, together with the sequence of feasible flight corridors, denoted by $\{\mathcal{P}_i^H = [\mathbf{G}_i^H, \mathbf{h}_i^H]^T\}_{i=1}^{M_{max}}$, to the sequence of optimal

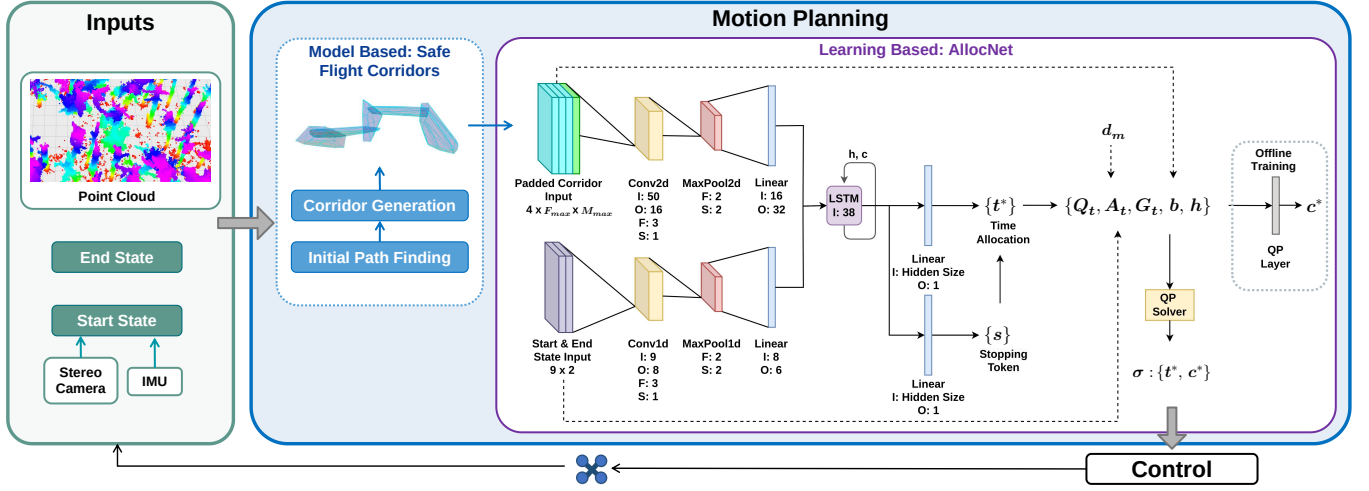


Fig. 3. The overall planning framework. The network architecture is included in the motion planning part. With some initial path generation and corridor generation, we can use two terms of network input: padded corridor and start and end state input. “I”, “O” represent the input and output size, “F” represents kernel size, and “S” is stride. The output layer for time allocation is connected with a Softplus function to ensure $t^* > 0$. During the offline training, we utilize a QP implicit layer to propagate gradients toward time allocation, and for online planning, we directly use a QP solver to solve the minimum control trajectory optimization.

times, $t^* \in \mathbb{R}^M$. Each flight corridor is represented as one matrix in $\mathbb{R}^{F_{max} \times 4}$. Here F_{max} represents the maximum number of faces of the corridor, that is, the number of linear inequality constraints that define it, and 4 corresponds to the encoding of each such inequality using the unit normal and the signed distance of the hyperplane that represents its zeroth level set from the origin. We note that the actual number of polytopes M need not be fixed; it can take on any value at most M_{max} . At the level of inputs to *AllocNet*, this is handled by zeroing out the data corresponding to the trailing P_i^H . We further introduce another, auxiliary, output of the network - a “stopping token” [33] that signals the maximum index of the M -dimensional output vector t that corresponds to the traversal time of a supplied corridor.

2) *Implicit differentiable layers*: With time allocation from the previous layer, we can explicitly formulate the trajectory optimization problem as a quadratic programming problem. The matrices formed by time allocation are defined as a set $\Phi(t) = \{Q(t), A(t), G(t)\}$. Then we introduce a differentiable layer with its input and output as:

$$z_{i+1} = \arg \min_z z^T Q(z_i) z, \quad (5a)$$

$$\text{s.t. } A(z_i) z = b, \quad (5b)$$

$$G(z_i) z \leq h, \quad (5c)$$

where $z_i = t^*$ refers to the optimal time allocation from the previous layer, $z_{i+1} = c^*$ is the optimal coefficient matrix. After solving this problem, we can get the equalities part of the KKT condition [34] with the primal-dual optimal solution,

$$\Gamma(c^*, \nu^*, \lambda^*, \Phi) = \begin{bmatrix} 2Qc^* + A^T \nu^* + G^T \lambda^* \\ Ac^* - b \\ \text{diag}(\lambda^*)(Gc^* - h) \end{bmatrix} = 0. \quad (6)$$

To keep it short, we use $y^* = (c^*, \nu^*, \lambda^*)$ in the rest of paper. By using implicit differentiation, we can get

$$\frac{\partial \Gamma(y^*, \Phi)}{\partial \Phi} + \frac{\partial \Gamma(y^*, \Phi)}{\partial y^*} \frac{\partial y^*}{\partial \Phi} = 0. \quad (7)$$

Therefore, we can expand the partial derivative of output towards input as time allocation with chain rules by:

$$\frac{\partial y^*}{\partial t^*} = \frac{\partial y^*}{\partial \Phi} \frac{\partial \Phi}{\partial t^*} = - \left(\frac{\partial \Gamma(y^*, \Phi)}{\partial y^*} \right)^{-1} \frac{\partial \Gamma(y^*, \Phi)}{\partial \Phi} \frac{\partial \Phi}{\partial t^*}. \quad (8)$$

With the above expression, we are able to backpropagate gradients of the primal solution c^* towards time t^* .

3) *Loss functions*: We consider two main loss functions, one loss term ℓ_F for min objective and a stopping token loss ℓ_T for a flexible number of segments. The loss function is evaluated by the objective of the original nonlinear programming

$$\ell_F(t^*) = (c^*)^T Q(t^*) c^* + w_t \mathbf{1}^T t^*, \quad (9)$$

where $\mathbf{1}$ is a vector of ones to minimize the total time term. We need to propagate the gradient toward the previous layer. Based on Eq. (8), we can further utilize chain rules to get all gradients to time allocation, as:

$$\nabla_t \ell_F(t^*) = \left(\frac{\partial Q}{\partial t^*} \right)^T \nabla_Q \ell_F + \left(\frac{\partial c^*}{\partial t^*} \right)^T \nabla_{c^*} \ell_F + w_t \mathbf{1} \quad (10)$$

The QP implicit layer assumes the feasibility of the problem and the satisfaction of the KKT condition. However, there are no guarantees to formulate problems that are always feasible primarily because of invalid time allocations. Additionally, in certain scenarios, the QP solver cannot find an optimal solution within the allowed time constraints. In differentiable optimization layers with learning hard constraints, usually, the relaxed problems are studied to ensure feasibility, which will leave the gap to the original problem [16]. For the conditions

of solver failures, we instead use a different loss term to penalize the deviation to the reference feasible time allocation,

$$\ell_F(\mathbf{t}^*) = w_F \|\bar{\mathbf{t}} - \mathbf{t}^*\|^2 + w_t \mathbf{1}^T \mathbf{t}^*, \quad (11)$$

where w_F is the weight for the MSE loss.

When generating the overlapping flight corridors, the number of segments usually depends on the initial seed path for polytopes and goal position, which we expect to represent the problem efficiently. To maintain such an efficient representation, a flexible segment length rather than a redundant corridor sequence would be much better for formulating the optimization. Thus, we design the stopping token loss for *AllocNet* that considers the timeliness of terminating the sequence.

More specifically, the token loss function aims to minimize three terms: the prediction error, the cost associated with prematurely ending the sequence, and the cost associated with ending the sequence late. The token loss function ℓ_T is formulated as follows:

$$\ell_T(\mathbf{s}) = \ell_{BCE}(\mathbf{s}, \bar{\mathbf{s}}) + \lambda_p \cdot (\ell_{PEP} + \ell_{LEP}), \quad (12)$$

where $\mathbf{s} = [s_1, \dots, s_{M_{max}}]$ represents the predicted stop tokens, $\bar{\mathbf{s}} = [\bar{s}_1, \dots, \bar{s}_{M_{max}}]$ is the ground truth stop tokens, ℓ_{BCE} is the Binary Cross Entropy loss, ℓ_{PEP} is the premature end penalty, ℓ_{LEP} is the late end penalty, and λ_p is the penalty weight for the premature and late end penalties. We can define the premature end penalty ℓ_{PEP} and late end penalty ℓ_{LEP} as:

$$\ell_{PEP} = \sum_{i=1}^{M_{max}} \mathbb{1}_{\{s_i > \alpha \wedge \bar{s}_i < \alpha\}}, \quad (13)$$

$$\ell_{LEP} = \sum_{i=1}^{M_{max}} \mathbb{1}_{\{s_i < \alpha \wedge \bar{s}_i > \alpha\}}, \quad (14)$$

where α is the token threshold, and $\mathbb{1}_{\{\cdot\}}$ is the indicator function which equals 1 when the condition inside the braces is true and 0 otherwise. For each stop token, 0 and 1 indicate the mid and the end of a sequence, respectively.

Combining ℓ_F and ℓ_T , *AllocNet* is therefore trained with the following cumulative loss function

$$\ell = \ell_F(\mathbf{t}^*) + w_S \cdot \ell_T(\mathbf{s}), \quad (15)$$

where w_S is the weight for balancing ℓ_F and ℓ_T .

IV. RESULTS

A. Numerical Evaluation

1) *Dataset generation*: We use the real-world flying point cloud dataset M3ED [35] and randomly generated maps with Gaussian noise toward obstacles in the simulation for training. To model local map regions for online planning, the point cloud data is sampled and cropped into the same size as $12.5m \times 12.5m \times 5m$, re-centered into the same origin. We sample 100 valid start and goal states in each map to get 260000 data groups. We apply Informed RRT* [36] to generate an initial geometric path and use this path to get convex covers as continuous overlapped corridors [1]. Because the runtime increases with the number of hyperplanes, we limit

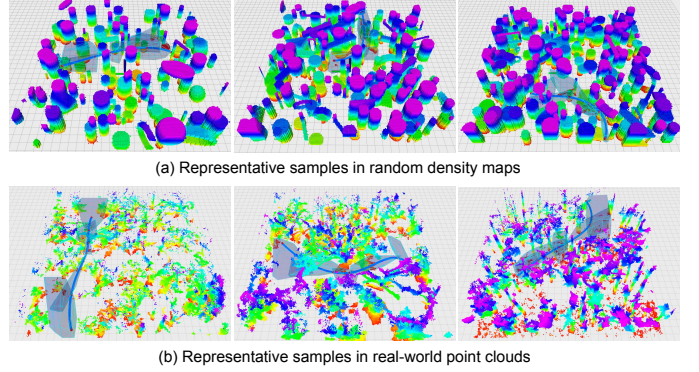


Fig. 4. Representative samples in simulation and real-world environments. The simulated map is randomly generated with different types of obstacles and different densities. The real-world flight data are cropped and re-centered. Some sampled paths and corridors are shown in the figure.

the maximum number of hyperplanes within each corridor to 50. Furthermore, we enforce a maximum trajectory segment of 5 or 10, or this data will be dropped. The training variables and computation time will increase with the number of segments based on the problem formulation. We use the well-known trajectory optimization methods [11] to get the reference feasible time allocation, which enables minimizing the non-fix total time. Some representative samples in both types of maps are shown in Fig. 4.

2) *Implementation Details*: We use 3-D 5 (or 7)-degree polynomials and solve the quadratic programming via OSQP [37]. The corridors and flat-based dynamic constraints are approximated discrete hard constraints for designed feasible trajectories. We implemented the *AllocNet* in PyTorch and trained it on our generated dataset with a carefully selected set of parameters, shown in Tab. I. The LSTM component within the network has a hidden size of 256 and comprises a single layer. The network is trained on NVIDIA RTX A5000 GPU and later integrated with other ROS modules by using LibTorch. The model can converge easily in 100 to 500 epochs during training with the reference time loss term.

TABLE I
PARAMETERS

v_{max}	a_{max}	j_{max}	N_{res}	w_F	w_t	w_S	λ_p
4 m/s	6 m/s ²	8m/s ³	20	1200	17.5	20	5

3) *Ablation Experiment*: We conduct a set of ablation experiments to comprehensively evaluate the performance of *AllocNet* and gain deeper insights into its mechanisms. These experiments involve comparing the learned *AllocNet* model with versions that include modifications to the loss function and part of the architecture. To evaluate the role and impact of the individual components and parameters of the proposed model, we prepare several alternative versions of the model:

- **Reference Time ℓ_F Only**: By training a model using only the ℓ_F loss function, we aim to investigate the role and contribution of ℓ_T to the performance of *AllocNet*.
- **Reference Time ℓ_F w/ Token**: In this variant, we set the token threshold α as 0.5, which is the most frequently

used token threshold for LSTM models in general. This experiment is designed to demonstrate the influence of tuned token thresholds on predicting the correct sequence length.

- **Ours w/ MLP Output Module:** We replace the LSTM output module with a 3-layer MLP with a hidden size of 256. This variant helps us evaluate the advantages of LSTM to predict segment numbers accurately.
- **Ours ($\alpha = 0.35/0.75$):** We evaluate the model with different stopping token threshold α to investigate the tradeoff of biased threshold to control the actual output length and the final performance.

To ensure attributable performance changes to specific modifications, all comparative models maintain the original parameter values for elements that have not been specifically altered for the experiment.

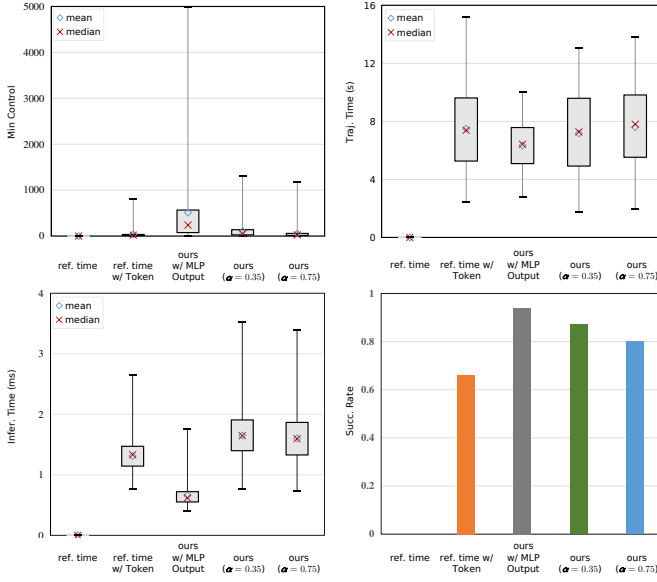


Fig. 5. Ablation study on minimum control efforts, total flight time (s), inference time (ms), and success rates.

The test result is shown in Fig. 5. With the use of only objective loss, the value of the output time vector tends to accumulate in the first index and barely works as expected. Adding the token loss significantly alleviates this issue but still stays in highly sub-optimal areas with a larger total time range and lower success rate. Although using an MLP layer instead of LSTM shows a higher success rate by consistently outputting the maximum length of values, it lacks precision in predicting the actual length of time sequences due to its non-autoregressive nature. In contrast, the autoregressive nature of LSTM, when trained with token loss, allows for more accurate termination of sequences, leading to lower control efforts compared to MLP, which tends to overfit the most frequent sequence length. For our proposed method, the failure scenarios are mainly because the length of valid time allocation doesn't fit the actually input corridor length, which is nontrivial to train together with multiple loss functions. The stopping token threshold is usually set as 0.5, but with multiple loss functions, relaxing the threshold ($\alpha = 0.35$) will

accordingly increase the overall success rate. Otherwise, using larger threshold results in more conservative total trajectory times and lower control efforts.

B. Benchmark Comparisons in Simulation

We compare our proposed methods with some existing solving schemes. All methods are provided with dynamic and corridor constraints, with the trajectory represented by piecewise polynomials. Theoretically, using linear inverse with higher derivatives waypoints equality constraints can achieve a more numerically stable solution [17], and our sub-problem formulation can also be replaced by this method. To make a fair comparison, we only choose benchmark methods that can be combined with general nonlinear or quadratic programming formulation, and all are formulated with the same general inequality constraints. We also observe some learning-based methods with hard constraints [30, 32], but the performances are not usually better than the state-of-the-art solver and are even worse for nonlinear problems. We evaluate all these methods in four different types of environments: simulated random map, forest, urban, and indoor. We utilize only the simulation and forest data during training, whereas the urban and indoor maps come from unseen data distributions. The details of benchmark methods are:

- **QP*** [3]: The quadratic programming is solved by OSQP, and the optimal times' gradient descent is solved by backtracking line search using L-BFGS [38]. The original problem doesn't enforce dynamic inequality constraints; therefore, it can reach a 100% success rate. However, we include all these constraints for comparison and will directly exit the line search if it reaches an invalid time allocation.
- **NLP**: we use NLOpt [39] to solve the nonlinear programming in Prob. 1 and use the Improved Stochastic Ranking Evolution Strategy (ISRES) algorithm to get the global optimal solution fast, we use the QP solution as an initial guess and limit 5 seconds for maximum solving time.
- **Bilevel** [14]: The time allocation and trajectory optimization are combined in a BiLevel Optimization problem. We replace the original trajectory representation as Bezier Curves to polynomials and also change the constraints formulation accordingly. The inner loop quadratic programming is solved by OSQP.

As indicated in Tab. II, we apply a traditional method to generate the safe flight corridor for all different datasets, ensuring there are no generalization issues in our proposed framework and comparison. In addition, we set the same relative tolerance criteria and limit the iterations for all the solvers. The NLP methods usually cannot converge with the allowed time, and the performance could be worse if no feasible initial guess is provided. The QP* and Bilevel methods need an initial total time, whereas other methods can directly optimize the time during optimization. The QP* method uses finite differences as the gradient direction and has a relatively lower success rate among all other methods. This is due to the fact that the quadratic programming will easily become infeasible

TABLE II
BENCHMARK ON TRAJECTORY OPTIMIZATION

	Method	Min Control	Traj. Time (s)	Comp. Time (ms)	Succ. Rate
Sim	QP*	129.08	7.66	1820.15	0.51
	NLP	154.61	7.00	5095.19	0.79
	Bilevel OPT	151.18	7.10	2030.81	0.76
	Ours ($\alpha = 0.35$)	108.32	7.00	55.36	0.84
	Ours ($\alpha = 0.75$)	55.14	7.54	44.91	0.80
Forest	QP*	108.05	7.89	1619.20	0.52
	NLP	138.26	7.22	5093.21	0.78
	Bilevel	136.79	7.29	1813.08	0.77
	Ours ($\alpha = 0.35$)	105.21	7.17	55.48	0.87
	Ours ($\alpha = 0.75$)	51.75	7.66	41.56	0.79
Urban	QP*	104.86	8.05	1594.43	0.51
	NLP	129.62	7.22	5095.45	0.81
	Bilevel	124.32	7.40	1835.72	0.79
	Ours ($\alpha = 0.35$)	101.03	7.23	54.39	0.86
	Ours ($\alpha = 0.75$)	54.97	7.63	43.29	0.80
Indoor	QP*	104.63	8.02	1481.62	0.52
	NLP	130.90	7.32	5089.54	0.80
	Bilevel	127.06	7.43	1769.22	0.78
	Ours ($\alpha = 0.35$)	102.64	7.34	52.23	0.90
	Ours ($\alpha = 0.75$)	56.09	7.74	42.31	0.79

because of invalid time allocation. The benchmark results highlight the advantages in terms of both computation time and also minimum control efforts of our proposed method. Even when the total solving time for constrained quadratic programming is factored in, our proposed method can still achieve fast computation time and a reasonable success rate for online planning. For online planning, we can combine multiple models to get better overall performance in practice.



Fig. 6. Snapshots of flight experiments in the outdoor (a) and indoor (b) cluttered environments.

C. Real-world Deployment

We validate the proposed method on our custom-designed Falcon 250 v2 platform [2], with ModalAI VOXL board¹ and Intel Realsense camera D435i for real-time state estimation and volumetric mapping. The inference runs around 1.0 ms on an i7-10710U CPU, which is similar to the computation time achieved when running on both CPU and GPU versions. Given the limited sensor range, the box-bounded maximum velocity and acceleration are set to 2.0 m/s to 3.0 m/s² respectively. We utilize a finite horizon local planning style with a 100 HZ finite state machine to call the trajectory optimization module. When calling the replanning with nonhomogeneous boundary conditions, we also employ models trained from non-rest to rest states models to handle these scenarios. Since we use simple geometrical pathfinding to generate corridors, the performance would be better if using some dynamic planning methods to get dynamically feasible corridors. We extensively demonstrate our framework in both indoor and outdoor cluttered environments, and the snapshots are shown in Fig. 6.

V. CONCLUSION

In this paper, we introduce an interpretable and lightweight learning-based planning framework that utilizes neural networks to learn the optimal time allocation for constrained trajectory optimization problems. We evaluate our proposed method through extensive simulations and indoor and outdoor experiments demonstrating its efficacy for real-time safe operation. Our proposed method relies on the efficiency of the quadratic programming solver and the quality of the initial path. The latter forces the solution to lie in the same homotopy class as the initial path for corridor generation. Our future work will focus on integrating efficient and optimal search-based path initialization and corridor generation on GPUs.

VI. ACKNOWLEDGMENTS

We would like to thank Fernando Cladera, Yuezhan Tao, Yifei (Simon) Shao, and Alex Zhou for helping with the results in the paper. The data sets used here are described in [35]. The experimental platform was developed for the work presented in [2].

REFERENCES

- [1] S. Liu, M. Watterson, K. Mohta, K. Sun, S. Bhattacharya, C. J. Taylor, and V. Kumar, "Planning dynamically feasible trajectories for quadrotors using safe flight corridors in 3-d complex environments," *IEEE Robotics and Automation Letters*, vol. 2, no. 3, pp. 1688–1695, 2017.
- [2] Y. Tao, Y. Wu, B. Li, F. Cladera, A. Zhou, D. Thakur, and V. Kumar, "Seer: Safe efficient exploration for aerial robots using learning to predict information gain," in *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 2023, pp. 1235–1241.
- [3] D. Mellinger and V. Kumar, "Minimum snap trajectory generation and control for quadrotors," in *2011 IEEE International Conference on Robotics and Automation*, 2011, pp. 2520–2525.
- [4] R. Penicka, Y. Song, E. Kaufmann, and D. Scaramuzza, "Learning minimum-time flight in cluttered environments," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 7209–7216, 2022.

¹<https://www.modalai.com/pages/voxl>

- [5] S. Ross, N. Melik-Barkhudarov, K. S. Shankar, A. Wendel, D. Dey, J. A. Bagnell, and M. Hebert, "Learning monocular reactive uav control in cluttered natural environments," in *2013 IEEE International Conference on Robotics and Automation*, 2013, pp. 1765–1772.
- [6] A. Loquercio, E. Kaufmann, R. Ranftl, M. Müller, V. Koltun, and D. Scaramuzza, "Learning high-speed flight in the wild," *Science Robotics*, vol. 6, no. 59, p. eabg5810, 2021. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.abg5810>
- [7] J. Chen, T. Liu, and S. Shen, "Online generation of collision-free trajectories for quadrotor flight in unknown cluttered environments," in *2016 IEEE International Conference on Robotics and Automation (ICRA)*, 2016, pp. 1476–1483.
- [8] C. Richter, A. Bry, and N. Roy, "Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments," in *Robotics Research: The 16th International Symposium ISRR*. Springer, 2016, pp. 649–666.
- [9] J. Tordesillas and J. P. How, "FASTER: Fast and safe trajectory planner for navigation in unknown environments," *IEEE Transactions on Robotics*, 2021.
- [10] I. Spasojevic, V. Murali, and S. Karaman, "Perception-aware time optimal path parameterization for quadrotors," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*, 2020, pp. 3213–3219.
- [11] Z. Wang, X. Zhou, C. Xu, and F. Gao, "Geometrically constrained trajectory optimization for multi-copters," *IEEE Transactions on Robotics*, vol. 38, no. 5, pp. 3259–3278, 2022.
- [12] A. Bry, C. Richter, A. Bachrach, and N. Roy, "Aggressive flight of fixed-wing and quadrotor aircraft in dense indoor environments," *The International Journal of Robotics Research*, vol. 34, no. 7, pp. 969–1002, 2015. [Online]. Available: <https://doi.org/10.1177/0278364914558129>
- [13] F. Gao, W. Wu, J. Pan, B. Zhou, and S. Shen, "Optimal time allocation for quadrotor trajectory generation," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 4715–4722.
- [14] W. Sun, G. Tang, and K. Hauser, "Fast uav trajectory optimization using bilevel optimization with analytical gradients," in *2020 American Control Conference (ACC)*, 2020, pp. 82–87.
- [15] P. Foehn, A. Romero, and D. Scaramuzza, "Time-optimal planning for quadrotor waypoint flight," *Science Robotics*, vol. 6, no. 56, p. eabh1221, 2021. [Online]. Available: <https://www.science.org/doi/abs/10.1126/scirobotics.abh1221>
- [16] B. Amos and J. Z. Kolter, "Optnet: Differentiable optimization as a layer in neural networks," in *Proceedings of the 34th International Conference on Machine Learning - Volume 70*, ser. ICML'17. JMLR.org, 2017, p. 136–145.
- [17] M. M. de Almeida and M. Akella, "New numerically stable solutions for minimum-snap quadcopter aggressive maneuvers," in *2017 American Control Conference (ACC)*, 2017, pp. 1322–1327.
- [18] D. Burke, A. Chapman, and I. Shames, "Generating minimum-snap quadrotor trajectories really fast," in *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2020, pp. 1487–1492.
- [19] M. Burri, H. Oleynikova, M. W. Achtelik, and R. Siegwart, "Real-time visual-inertial mapping, re-localization and planning onboard mavs in unknown environments," in *2015 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2015, pp. 1872–1878.
- [20] S. Liu, N. Atanasov, K. Mohta, and V. Kumar, "Search-based motion planning for quadrotors using linear quadratic minimum time control," in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2017, pp. 2872–2879.
- [21] R. Penicka and D. Scaramuzza, "Minimum-time quadrotor waypoint flight in cluttered environments," *IEEE Robotics and Automation Letters*, vol. 7, no. 2, pp. 5719–5726, 2022.
- [22] J. Bobrow, S. Dubowsky, and J. Gibson, "Time-optimal control of robotic manipulators along specified paths," *The International Journal of Robotics Research*, vol. 4, no. 3, pp. 3–17, 1985. [Online]. Available: <https://doi.org/10.1177/027836498500400301>
- [23] S. Bansal, V. Tolani, S. Gupta, J. Malik, and C. Tomlin, "Combining optimal control and learning for visual navigation in novel environments," in *Proceedings of the Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, L. P. Kaelbling, D. Kragic, and K. Sugiura, Eds., vol. 100. PMLR, 30 Oct–01 Nov 2020, pp. 420–429. [Online]. Available: <https://proceedings.mlr.press/v100/bansal20a.html>
- [24] A. Faust, O. Ramirez, M. Fiser, K. Oslund, A. Francis, J. Davidson, and L. Tapia, "Prm-rl: Long-range robotic navigation tasks by combining reinforcement learning and sampling-based planning," in *IEEE International Conference on Robotics and Automation (ICRA)*, Brisbane, Australia, 2018, pp. 5113–5120. [Online]. Available: <https://arxiv.org/abs/1710.03937>
- [25] G. Tang, W. Sun, and K. Hauser, "Learning trajectories for real-time optimal control of quadrotors," in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018, pp. 3620–3625.
- [26] F. Yang, C. Wang, C. Cadena, and M. Hutter, "iplanner: Imperative path planning," in *Robotics: Science and Systems (RSS)*, 2023. [Online]. Available: <https://arxiv.org/pdf/2302.11434.pdf>
- [27] G. Ryou, E. Tal, and S. Karaman, "Real-time generation of time-optimal quadrotor trajectories with semi-supervised seq2seq learning," in *Proceedings of The 6th Conference on Robot Learning*, ser. Proceedings of Machine Learning Research, K. Liu, D. Kulic, and J. Ichnowski, Eds., vol. 205. PMLR, 14–18 Dec 2023, pp. 1860–1870. [Online]. Available: <https://proceedings.mlr.press/v205/ryou23a.html>
- [28] A. Agrawal, B. Amos, S. Barratt, S. Boyd, S. Diamond, and J. Z. Kolter, "Differentiable convex optimization layers," *Advances in neural information processing systems*, vol. 32, 2019.
- [29] N. Jaquier, Y. Zhou, J. Starke, and T. Asfour, "Learning to sequence and blend robot skills via differentiable optimization," *IEEE Robotics and Automation Letters*, vol. 7, no. 3, pp. 8431–8438, 2022.
- [30] P. Donti, D. Rolnick, and J. Z. Kolter, "Dc3: A learning method for optimization with hard constraints," in *International Conference on Learning Representations*, 2021.
- [31] G. Négier, M. W. Mahoney, and A. Krishnapriyan, "Learning differentiable solvers for systems with hard constraints," in *The Eleventh International Conference on Learning Representations*, 2023. [Online]. Available: <https://openreview.net/forum?id=vdv6CmGksr0>
- [32] T. Frerix, M. Nießner, and D. Cremers, "Homogeneous linear inequality constraints for neural network activations," in *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2020, pp. 3229–3234.
- [33] I. Sutskever, O. Vinyals, and Q. V. Le, "Sequence to sequence learning with neural networks," in *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, ser. NIPS'14. Cambridge, MA, USA: MIT Press, 2014, p. 3104–3112.
- [34] S. Boyd and L. Vandenberghe, *Convex optimization*. Cambridge university press, 2004.
- [35] K. Chaney, F. Cladera, Z. Wang, A. Bisulco, M. A. Hsieh, C. Korpela, V. Kumar, C. J. Taylor, and K. Daniilidis, "M3ed: Multi-robot, multi-sensor, multi-environment event dataset," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*, June 2023, pp. 4015–4022.
- [36] J. D. Gammell, S. S. Srinivasa, and T. D. Barfoot, "Informed rrt*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic," in *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2014, pp. 2997–3004.
- [37] B. Stellato, G. Banjac, P. Goulart, A. Bemporad, and S. Boyd, "OSQP: an operator splitting solver for quadratic programs," *Mathematical Programming Computation*, vol. 12, no. 4, pp. 637–672, 2020. [Online]. Available: <https://doi.org/10.1007/s12532-020-00179-2>
- [38] D. C. Liu and J. Nocedal, "On the limited memory bfgs method for large scale optimization," *Mathematical programming*, vol. 45, no. 1-3, pp. 503–528, 1989.
- [39] S. G. Johnson, "The NLOpt nonlinear-optimization package," <https://github.com/stevengi/nlopt>, 2007.