

NAVIGATION

Learning high-speed flight in the wild

Antonio Loquercio^{1*†}, Elia Kaufmann^{1†}, René Ranftl², Matthias Müller², Vladlen Koltun³, Davide Scaramuzza¹

Quadrotors are agile. Unlike most other machines, they can traverse extremely complex environments at high speeds. To date, only expert human pilots have been able to fully exploit their capabilities. Autonomous operation with onboard sensing and computation has been limited to low speeds. State-of-the-art methods generally separate the navigation problem into subtasks: sensing, mapping, and planning. Although this approach has proven successful at low speeds, the separation it builds upon can be problematic for high-speed navigation in cluttered environments. The subtasks are executed sequentially, leading to increased processing latency and a compounding of errors through the pipeline. Here, we propose an end-to-end approach that can autonomously fly quadrotors through complex natural and human-made environments at high speeds with purely onboard sensing and computation. The key principle is to directly map noisy sensory observations to collision-free trajectories in a receding-horizon fashion. This direct mapping drastically reduces processing latency and increases robustness to noisy and incomplete perception. The sensorimotor mapping is performed by a convolutional network that is trained exclusively in simulation via privileged learning: imitating an expert with access to privileged information. By simulating realistic sensor noise, our approach achieves zero-shot transfer from simulation to challenging real-world environments that were never experienced during training: dense forests, snow-covered terrain, derailed trains, and collapsed buildings. Our work demonstrates that end-to-end policies trained in simulation enable high-speed autonomous flight through challenging environments, outperforming traditional obstacle avoidance pipelines.

INTRODUCTION

Quadrotors are among the most agile and dynamic machines ever created (1, 2). Thanks to their agility, they can traverse complex environments, ranging from cluttered forests to urban canyons, and reach locations that are otherwise inaccessible to humans and machines alike. This ability has led to their application in fields such as search and rescue, logistics, security, infrastructure, entertainment, and agriculture (3). In most of these existing applications, the quadrotor needs to be controlled by expert human pilots, who take years to train and are thus an expensive and scarce resource. Infusing quadrotors with autonomy—that is, the capability to safely operate in the world without the need for human intervention—has the potential to massively enhance their usefulness and to revolutionize whole industries. However, the development of autonomous quadrotors that can navigate in complex environments with the agility and safety of expert human pilots or birds is a long-standing challenge that is still open.

The limiting factor for autonomous agile flight in arbitrary unknown environments is the coupling of fast and robust perception with effective planning. The perception system has to be robust to disturbances such as sensor noise, motion blur, and changing illumination conditions. In addition, an effective planner is necessary to find a path that is both dynamically feasible and collision free while relying only on noisy and partial observations of the environment. These requirements, together with the limited computational resources that are available onboard, make it difficult to achieve reliable perception and planning at low latency and high speeds (4).

Various approaches to enable autonomous flight have been proposed in the literature. Some works tackle only perception and build

high-quality maps from imperfect measurements (5–9), whereas others focus on planning without considering perception errors (10–13). Numerous systems that combine online mapping with traditional planning algorithms have been proposed to achieve autonomous flight in previously unknown environments (14–22). A taxonomy of prior works is presented in fig. S5 in the Supplementary Materials.

The division of the navigation task into the mapping and planning subtasks is attractive from an engineering perspective because it enables parallel progress on each component and makes the overall system interpretable. However, it leads to pipelines that largely neglect interactions between the different stages and thus compound errors (22). Their sequential nature also introduces additional latency, making high-speed and agile maneuvers difficult to impossible (4). Although these issues can be mitigated to some degree by careful hand-tuning and engineering, the divide-and-conquer principle that has been prevalent in research on autonomous flight in unknown environments for many years imposes fundamental limits on the speed and agility that a robotic system can achieve (23).

In contrast to these traditional pipelines, some recent works propose to learn end-to-end policies directly from data without explicit mapping and planning stages (24–27). These policies are trained by imitating a human (24, 27) from experience that was collected in simulation (25) or directly in the real world (26). Because the number of samples required to train general navigation policies is very high, existing approaches impose constraints on the quadrotor's motion model, for example, by constraining the platform to planar motion (24, 26, 27) and/or discrete actions (25), at the cost of reduced maneuverability and agility. More recent work has demonstrated that very agile control policies can be trained in simulation (28). Policies produced by the last approach can successfully perform acrobatic maneuvers but can only operate in unobstructed free space and are essentially blind to obstacles in the environment.

Copyright © 2021
The Authors, some
rights reserved;
exclusive licensee
American Association
for the Advancement
of Science. No claim
to original U.S.
Government Works

Downloaded from https://www.science.org at The Hong Kong University of Science and Technology (Guangzhou) on May 26, 2026

¹University of Zurich, Zürich, Switzerland. ²Intel, Munich, Germany. ³Intel, Santa Clara, CA, USA.

*Corresponding author. Email: loquercio@ifi.uzh.ch

†These authors contributed equally to this work.

Here, we present an approach to fly a quadrotor at high speeds in a variety of environments with complex obstacle geometry (Fig. 1 and Movie 1) while having access to only onboard sensing and computation. By predicting navigation commands directly from sensor measurements, we decrease the latency between perception and action while simultaneously being robust to perception artifacts, such as motion blur, missing data, and sensor noise. To deal with sample complexity and not endanger the physical platform, we train the policy exclusively in simulation. We leverage abstraction of the input data to transfer the policy from simulation to reality (28, 29). To this

end, we use a stereo-matching algorithm to provide depth images as input to the policy. We show that this representation is both rich enough to safely navigate through complex environments and abstract enough to bridge simulation and reality. Our choice of input representation guarantees a strong similarity of the noise models between simulated and real observations and gives our policy robustness against common perceptual artifacts in existing depth sensors.

We train the navigation policy via privileged learning (30) on demonstrations that are provided by a sampling-based expert. Our expert has privileged access to a representation of the environment in the form of a three-dimensional (3D) point cloud as well as perfect knowledge about the state of the quadrotor. Because simulation does not impose real-time constraints, the expert additionally has an unconstrained computational budget. Whereas existing global planning algorithms (10, 12, 13) generally output a single trajectory, our expert uses Metropolis-Hastings (M-H) sampling to compute a distribution of collision-free trajectories. This captures the multimodal nature of the navigation task where many equally valid solutions can exist (for example, going either left or right around an obstacle). We therefore use our planner to compute trajectories with a short time horizon to ensure that they are predictable from onboard sensors and that the sampler remains computationally tractable. We bias the sampler toward obstacle-free regions by conditioning it on trajectories from a classic global planning algorithm (13).

We also reflect the multimodal nature of the problem in the design and training of the neural network policy. Our policy takes a noisy depth image and inertial measurements as sensory inputs and produces a set of short-term trajectories together with an estimate of individual trajectory costs. The trajectories are represented as high-order polynomials to ensure dynamical feasibility. We train the policy using a multihypothesis winner-takes-all loss that adaptively maps the predicted trajectories to the best trajectories that have been found by the sampling-based expert. At test time, we use the predicted trajectory costs to decide which trajectory to execute in a receding horizon. The policy network is designed to be extremely lightweight, which ensures that it can be executed onboard the quadrotor at the update rates required for high-speed flight.

The resulting policy can fly a physical quadrotor in natural and human-made




Fig. 1. Deployment of the presented approach in challenging environments. Time-lapse illustrations of agile navigation at speeds between 5 and 10 m/s in a variety of environments (top to bottom): a mountain trail, a forest, an airplane hangar, and a disaster zone. The middle rows show the same environments, but from different viewpoints. The last row shows navigation through a collapsed building and above a derailed train. Please watch Movie 1 to get a better sense of the speed and agility of our approach.



Movie 1. Demonstrations of agile flight in challenging environments with purely onboard sensing and computation. The video also summarizes the training methodology of the navigation policy.

environments at speeds that are unreachable by existing methods. We achieve this in a zero-shot generalization setting: We train on randomly generated obstacle courses composed of simple off-the-shelf objects, such as schematic trees and a small set of convex shapes such as cylinders and cubes. We then directly deploy the policy in the physical world without any adaptation or fine-tuning. Our platform experiences conditions at test time that were never seen during training. Examples include high dynamic range (when flying from indoor environments to outdoor environments), poorly textured surfaces (indoor environments and snow-covered terrain), thick vegetation in forests, and the irregular and complex layout of a disaster scenario (Fig. 2). These results suggest that our methodology enables a multitude of applications that rely on agile autonomous drones with purely onboard sensing and computation.

RESULTS

Our experiments in simulation show that the proposed approach reduces the failure rate up to 10 times with respect to state-of-the-art methods. We confirm our results in a variety of real-world environments using a custom-built physical quadrotor; we deploy our policy trained in simulation without any further adaptations. In all experiments, the drone was provided with a reference trajectory, which is not collision free (Fig. 3C, depicted in red), to encode the intended flight path. This reference can be provided by a user or a higher-level planning algorithm. The drone is tasked to follow that flight path and make adjustments as necessary to avoid obstacles. Recordings of the experiments can be found in Movie 1.

High-speed flight in the wild

We tested our approach in diverse real-world environments, as illustrated in Figs. 1 and 2. High-speed flight in these environments is very challenging because of their complex structure (e.g., thick vegetation, thin branches, or collapsed walls) and multiple options available to avoid obstacles. In addition, a high-level understanding of the environment is necessary, for example, to pass through far-away openings (Fig. 3C) or nearby obstacles (Fig. 2M). Flying at high speeds also necessitates low-latency perception and robustness to sensor noise, which is worsened by challenging illumination conditions and low-texture surfaces (e.g., due to snow). At the same time, only limited computational resources are available onboard. Despite all of these challenges, our approach was able to successfully navigate in all environments that it was tested in. Note that our policy was trained in simulation and was never exposed to any of these environments or conditions at training time.

We measure performance according to success rate, i.e., the percentage of successful runs over the total number of runs, where we consider a run successful if the drone reaches the goal location within a radius of 5 m without crashing. We performed a total of 56 experiments at different speeds. Overall, our approach was always able to achieve the goal when not colliding into an obstacle. We report the cumulative and individual success rates at various speeds in Fig. 3 (D and E). Our experimental platform performs state estimation and depth perception onboard using vision-based sensors. A detailed description of the platform is available in section S3. We group the environments that were used for experiments into two classes, natural and human-made, and highlight the distinct

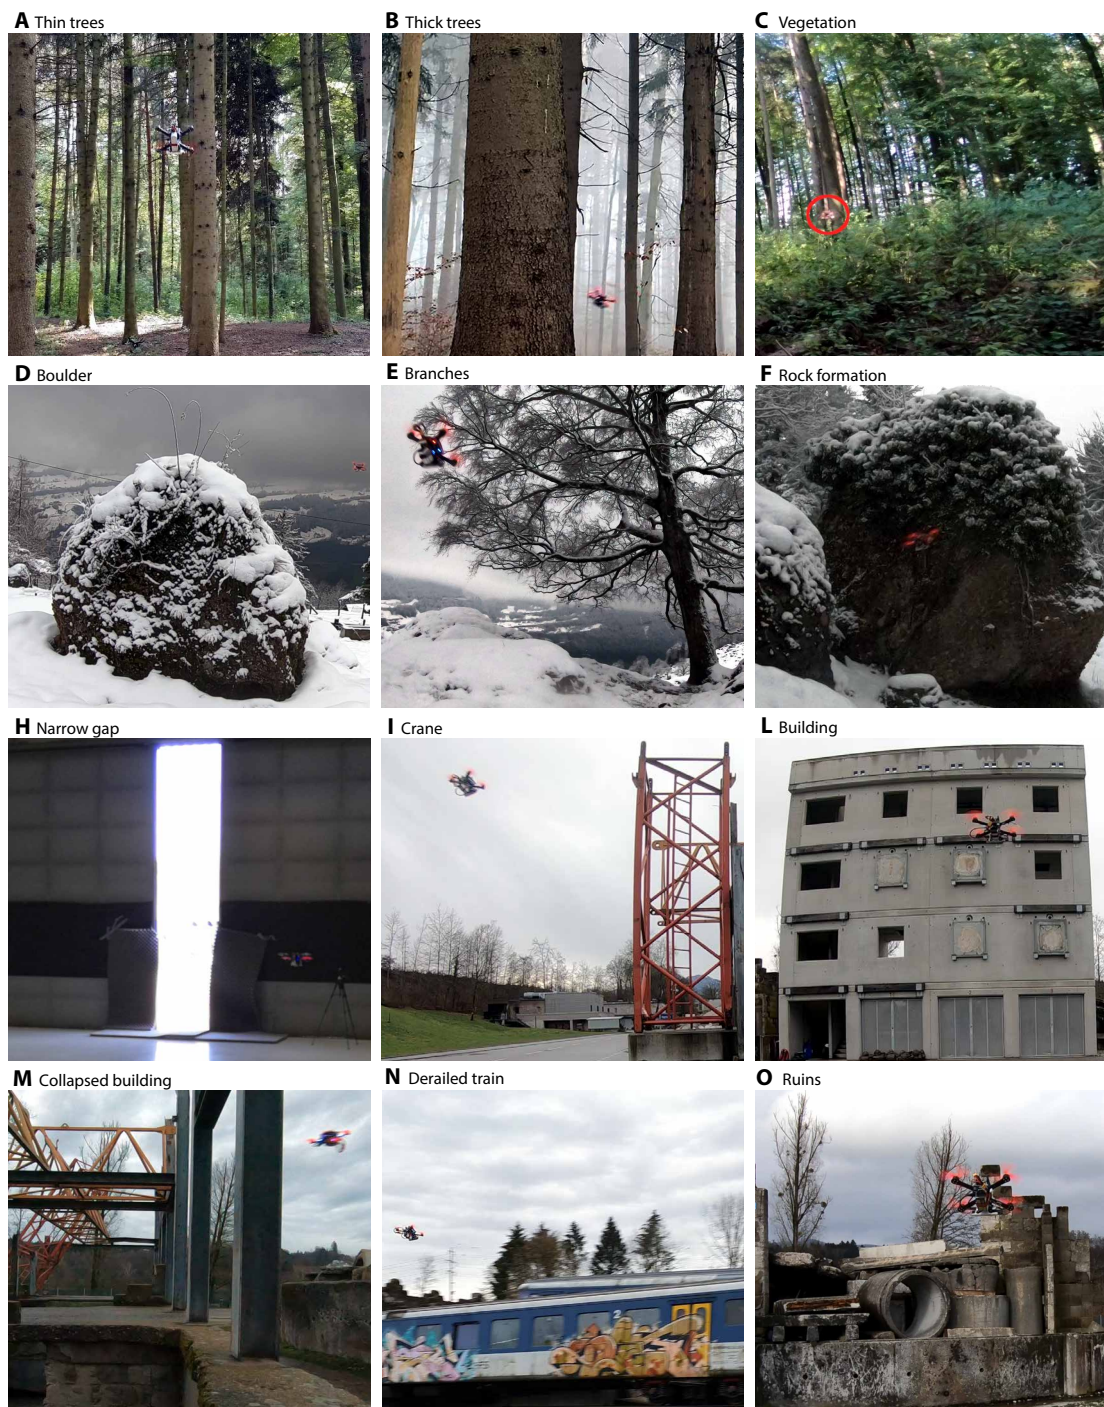


Fig. 2. Testing environments. Zero-shot generalization of our approach in complex natural (A to F) and human-made (H to O) environments. The encountered obstacles can often be avoided in multiple directions and have very different size and structure.

challenges that these types of environments present for agile autonomous flight.

Natural environments

We performed experiments in diverse natural environments: forests of different types and densities and steep snowy mountain terrains. Figure 2 (A to F) illustrates the heterogeneity of those environments. We performed experiments with two different reference trajectories:

a 40-m-long straight line and a circle with a 6-m radius (Fig. 3A). Both trajectory types are not collision free and would lead to a crash into obstacles if blindly executed. We flew the straight line at different average speeds in the range of 3 to 10 m/s. Flying at these speeds requires very precise and low-latency control of position and attitude to avoid bushes and pass through the openings between trees and branches (Fig. 3B). Traditional mapping and planning approaches

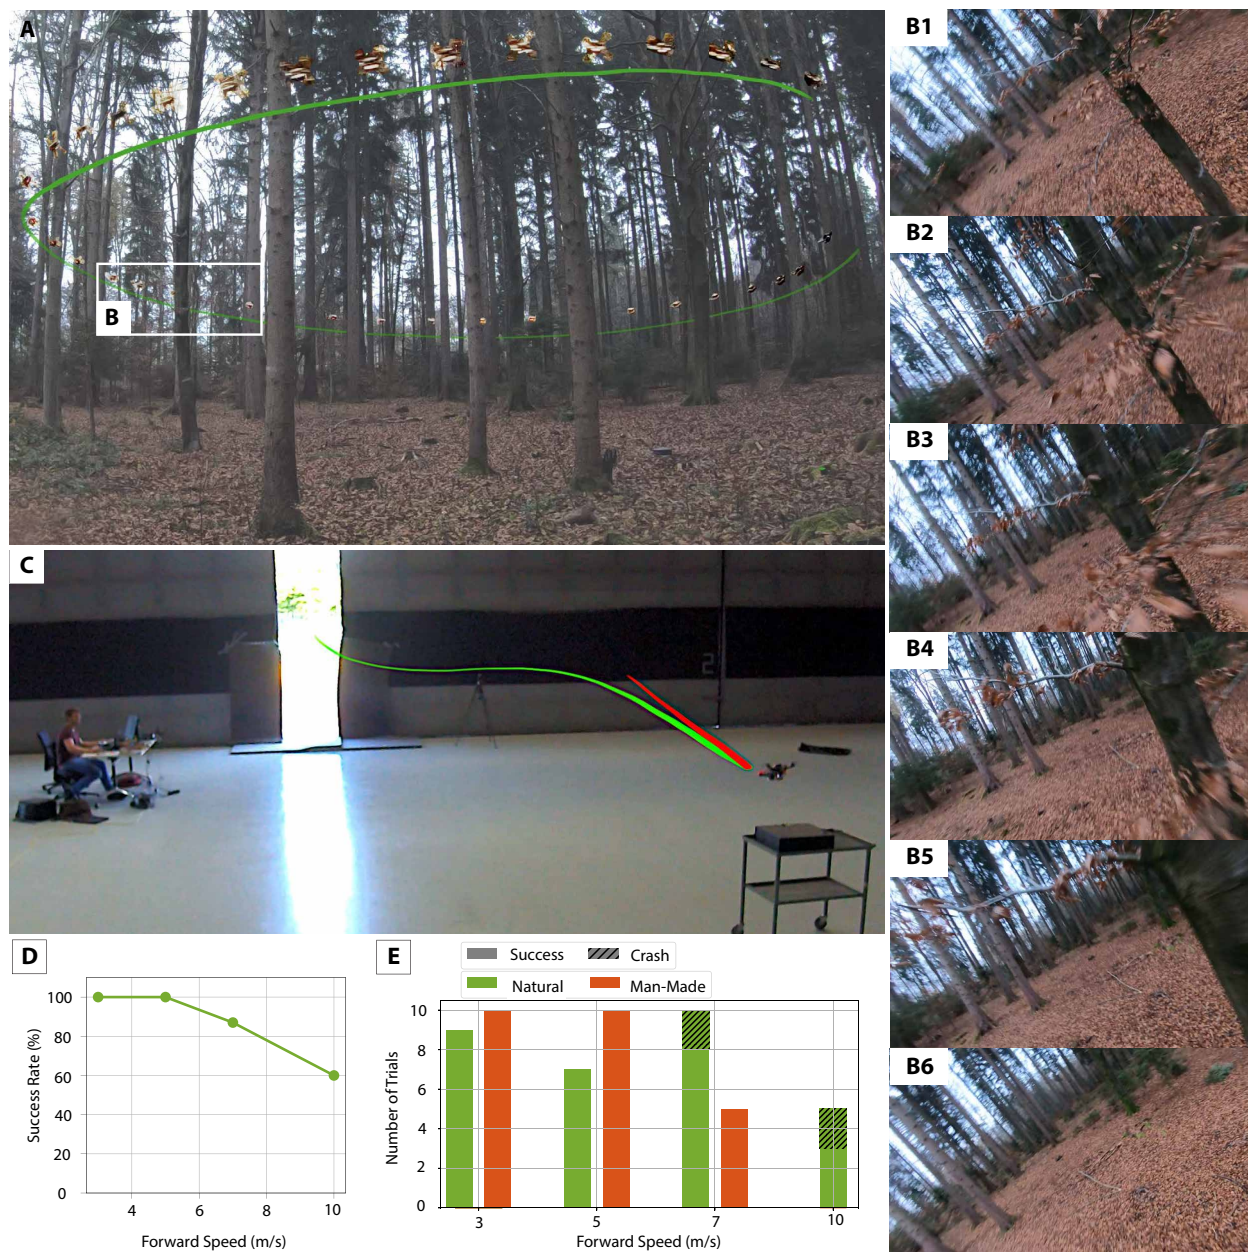


Fig. 3. Evaluation in indoor and outdoor environments. (A) Circular path in the forest at an average speed of 7 m/s. (B) Sequence of first-person views from the maneuver in (A) observed during the avoidance of tree branches. The maneuver requires fine and fast adaptations in height (B2 and B5) and attitude (B3 and B4) to avoid the vegetation. After the obstacle is passed, the drone accelerates in the direction of travel (B6). (C) Comparison of reference trajectory passing through a wall (in red) to actual flight path (in green) in an airplane hangar. (D) Success rates for all experiments aggregated according to flight speed. (E) Number of trials per environment class at different speeds.

generally fail to achieve high speeds in such conditions because both the thick vegetation and the texture-less snow terrain often cause very noisy depth observations. We conducted a total of 31 experiments in natural environments. At average speeds of 3 and 5 m/s, our approach consistently completed the task without experiencing a single crash. For comparison, state-of-the-art methods with comparable actuation, sensing, and computation (18, 31) achieve in similar environments a maximum average speed of 2.29 m/s. To study the limit of the system, we set the platform's average speed to 7 m/s. Despite the very high speed, the maneuver was successfully

completed in 8 of 10 experiments. The two failures happened when objects entered the field of view very late because of the high angular velocity of the platform. Given the good performance at 7 m/s, we push the average flight speed even higher to 10 m/s. At this speed, external disturbances, e.g., aerodynamics, battery power drops, and motion blur, start to play an important role and widen the simulation to reality gap. Nonetheless, we achieve a success rate of 60% with failures mainly happening in the proximity of narrow openings less than a meter wide, where a single wrong action results in a crash.

Human-made environments

We also tested our approach in a set of human-made environments, illustrated in Fig. 2 (G to O). In these environments, the drone faces a different set of challenges. It has to avoid obstacles with a variety of sizes and shapes (e.g., a train, a crane, a building, and ruins), slalom through concrete structures (c.f. Fig. 2M), and exit a building through a single narrow opening (c.f. Fig. 2H). The irregular and/or large structure of the encountered obstacles, the limited number of flyable openings, and the requirement to initiate the avoidance maneuver well in advance offer a complementary set of challenges with respect to our natural testing environments.

As in the natural environments, we provide the drone with a straight reference trajectory with length of 40 m. The reference trajectory is in direct collision with obstacles, and its blind execution would result in a crash. We performed a total of 19 experiments with flight speeds in the range of 3 to 7 m/s. Given the lower success rate experienced in the forest environment at 10 m/s, we did not test at higher speeds to avoid fatal crashes. As shown in Fig. 3E, our approach is robust to zero-shot deployment in these environments, whose characteristics were never observed at training time and consistently completes the task without crashing.

We further compare our approach to a commercial drone. The commercial drone (a Skydio R1) was tasked to follow a person running through the narrow gap. The speed of this drone cannot be enforced nor controlled and was therefore estimated in postprocessing. Specifically, the drones are required to exit a hangar by passing through a narrow gap of about 0.8 m in width (Fig. 3C). At the start of the experiment, the drones are placed at about 10 m in front of and about 5 m to the right of the gap. The task was represented by a straight reference trajectory passing through the wall (Fig. 3C, in red). This experiment is challenging because it requires a high-level understanding of the environment to turn early enough toward the gap. The commercial drone, flying at a speed of about 2.7 m/s, consistently failed to pass through the gap across three experiments. In two of the experiments, it deemed the gap to be too small and stopped in front of it; in the third, it crashed into the wall. In contrast, the low latency and robustness to perception failures of our approach enabled the drone to successfully fly through the narrow gap every time. We performed a total of six experiments at flight speeds of both 3 and 5 m/s and never experienced a crash.

Controlled experiments

We perform a set of controlled experiments in simulation to compare the performance of our approach with several baselines. We select two representative state-of-the-art approaches as baselines for navigation in unknown environments: the mapping and planning method of Zhou *et al.* (18) (FastPlanner) and the reactive planner of Florence *et al.* (32) (Reactive). The first approach (18) incrementally builds a map from a series of observations and plans a trajectory in this map to reach the goal while avoiding obstacles. In contrast, the second approach (32) does not build a map but uses instantaneous depth information to select the best trajectory from a set of predefined motion primitives based on a cost that encodes collision and progress toward the goal. The baselines receive the same input as our policy: the state of the platform, depth measurements from the stereo camera, and a goal in the form of a reference state that lies 5 s in the future. To provide a notion of the difficulty of the environments, we additionally show a naive baseline that blindly follows the reference trajectory to the goal without avoiding obstacles

(Blind). As in the real-world experiments, we compare the different approaches according to their success rate, which measures how often the drone reaches the goal location within a radius of 5 m without crashing. We perform all experiments in the Flightmare simulator (33) using the RotorS (34) Gazebo plugin for accurate physics modeling and Unity as a rendering engine (35). The experiments are conducted in four different environments: a forest, a narrow gap, a disaster scenario, and a city street. Those environments offer different challenges: from high-density but simple obstacle geometry in the forest to complex obstacle geometries in the urban scenarios. Sample observations collected from these environments are available in fig. S4 and movie S1. These environments are available out-of-the-box from the Unity Asset Store (<https://assetstore.unity.com/>).

The results of these experiments are summarized in Fig. 4. Throughout all environments, a similar pattern can be observed. At low speeds (3 m/s), all methods perform similarly. However, as the speed increases, the baselines' performances quickly drop: Already at 5 m/s, no baseline is able to complete all runs without crashing. In contrast, our method can reliably fly at high speeds through all environments, achieving an average success rate of 70% at 10. The drop in performance of the two baselines can be justified as follows. Although the reactive baseline has a low processing latency, it has a limited expressive power: The trajectory to be executed can only be selected from a relatively small set of primitives. This set is usually too small to account for all possible obstacle configurations and relative speeds between obstacles and drone, which are observed during high-speed flight. In addition, being the reactive baseline only conditioned on the current observation, it is strongly affected by noise in the observation. In contrast, the FastPlanner baseline can reject outliers in the depth map by leveraging multiple observations, which makes it more robust to sensing errors. However, this methods generally results in higher processing latency: Multiple observations are required to add obstacles in the map and, therefore, to plan trajectories to avoid them. This problem is worsened by high-speed motion, which generally results in little overlap between consecutive observations. By contrast, our approach has very the lowest processing latency (cf. the "Computational cost" section). In addition, the data-driven approach allows leveraging regularities in the data, which makes it more robust to sensor noise.

We thoroughly analyze those aspects in controlled studies on latency and sensor noise, which are shown in the effect of latency and sensor noise section. Additionally, study the robustness of our method to noise in estimation and control in section S1. A video demonstrating the performance of our method in the simulation environments can be found in movie S2. We now describe the specific characteristics of all environments and their respective experimental results in detail.

Forest

We build a simulated forest (36) in a rectangular region $R(l, w)$ of width $w = 30$ m and length $l = 60$ m, and fill it with trees that have a diameter of about 0.6 m. Trees are randomly placed according to a homogeneous Poisson point process P with intensity $\delta = 1/25$ tree/m² (36). As it is possible to see in Fig. 4, the resulting forest is very dense. We provide the drone with a straight reference trajectory of 40-m length and vary the average forward speed of the drone between 3 and 10 m/s. We repeat the experiments with 10 different random realizations of the forest using the same random seed for all baselines. At a low speed of 3 m/s, all methods successfully complete every run.




Fig. 4. Experiments in different simulated environments. **Left column:** Success rates at various speeds. **Right column:** Point cloud of the environment together with paths taken by different policies from start (green arrow) to end (red arrow). The paths taken by the blind and our policy are illustrated by red and green lines, respectively. Our approach consistently outperforms the baselines in all environments and at all speeds. Sample observations from these environments are shown in fig. S4.

As speed increases, the success rates of the baselines quickly degrade. At 10 m/s, no baseline completes even a single run successfully. In contrast, our approach is very robust at higher speeds. It achieves 100% success rate up to 5 m/s. At the highest speed of 10 m/s, our approach still keeps a success rate of 60%. An additional study of performance with respect to tree density δ is available in the Supplementary Materials (section S2).

Narrow gap

We then mimic in simulation the real-world narrow-gap experiment. We render a 40-m-long wall with a single opening, 10 m in front of the drone. The gap is placed at a randomized lateral offset in the range of $[-5, 5]$ m with respect to the starting location. The width of the opening is also randomized uniformly between 0.8 and 1.0 m. All experiments are repeated 10 times with different opening sizes

and lateral gap offsets for each speed. An illustration of the setup and the results of the evaluation are shown in Fig. 4. The reactive and blind baselines consistently fail to solve this task, whereas FastPlanner has a success rate of up to 50% at 5 m/s but still consistently fails at 7 m/s. We observe that the baselines adapt their trajectory only when being close to the wall, which is often too late to correct course toward the opening. Conversely, our approach always completes the task successfully at speeds of up to 5 m/s. Even at a speed of 7 m/s, it only fails in 2 of 10 runs.

Disaster and urban scenario

We further test our method in a simulated disaster zone (<https://youtu.be/z8ZtueSePIE>) and an urban city street (<https://youtu.be/gathQQfHEVo>). These environments are challenging because they contain obstacles with complex geometry. For instance, the former contains, among others, collapsed buildings, crashed cars, and iron fences, whereas the latter features street lamp poles, parked vehicles, and traffic signs. Some observations from these environments are shown in fig. S4. We start the drone from a random location in the map and provide it with a straight line reference of 40 m in length. For each starting location, the reference is guaranteed to be in collision with at least one obstacle. We vary the average flight speed from 3 to 10 m/s. The results of these experiments as well as the point clouds from one of the starting locations are shown in Fig. 4. Because of the complex obstacle geometries, even at a low speed of 3 m/s, the mapping baseline does not achieve 100% success rate. The reactive baseline performs better overall with a success rate of up to 90% at 5 m/s. However, its performance quickly degrades at higher speeds with as much as 8 of 10 failures at 10 m/s. Similarly to the narrow-gap experiments, we observe that the baselines only adapt their motion when very close to obstacles, which is often too late to avoid collision in the presence of large obstacles with complex shapes. In contrast to the baselines, our approach is much more robust at higher speeds with a success rate of up to 100% at 7 m/s. Even at a speed of 10 m/s, it only fails in 2 of 10 runs in the urban environment. The ability of our approach to combine long-term and

short-term planning is crucial to achieving this performance because the drone must steer early enough toward free space and at the same time perform small reactive corrections to avoid a collision. This ability, in addition to the low latency and robustness to sensor noise, gives our approach a substantial performance advantage with respect to the baselines in these challenging environments.

Computational cost

In this section, we compare the computational cost of our algorithm with the baselines. Table 1 shows the results of this evaluation. It highlights how each step of the methods contributes to the overall processing latency. All timings were recorded on a desktop computer with a 6-core i7-8700 central processing unit (CPU), which was also used to run the simulation experiments. To ensure a fair comparison, we report the timings when using only the CPU for all approaches. We also report the timings of our approach when performing neural network inference on a GeForce RTX 2080 graphics processing unit (GPU), because accelerators can be used with our approach without any extra effort. To paint a complete and realistic picture, we additionally evaluate the timing of our algorithm on the onboard computer of the quadrotor, which is a Jetson TX2.

With a total computation time of 65.2 ms per frame, FastPlanner incurs the highest processing latency. It is important to note that the temporal filtering operations that are necessary to cope with sensing errors effectively make perception even slower. Two to three observations of an obstacle can be required to add it to the map, which increases the effective latency of the system. By foregoing the mapping stage together, the Reactive baseline substantially reduces computation time. This baseline is about three times faster than FastPlanner with a total processing latency of 19.1 ms. However, the reduced processing latency comes at the cost of the trajectory complexity that can be represented because the planner can only select primitives from a predefined library. In addition, the reactive baseline is sensitive to sensing errors, which can drastically affect performance at high speeds.

Table 1. Processing latency (μ) on a desktop computer equipped with a hexacore i7-8700 CPU and a GeForce RTX 2080 GPU. The SD σ is computed over 1000 samples. For our approach, we report the computation time on the CPU and GPU (marked with *) on the desktop computer (Ours), as well as the computation time on the onboard computation unit Jetson TX2 (Ours Onboard). For the FastPlanner (18) and Reactive (32) baselines, preprocessing represents the time to build a point cloud from the depth image, whereas for our method, preprocessing is the time to convert depth into an input tensor for the neural network. More details on the subtasks division are available in the Supplementary Materials (section S4). The proposed methodology is substantially faster than prior works. NN, neural network.

Method	Components	μ (ms)	σ (ms)	Perception (%)	Total processing latency (ms)
FastPlanner (18)	Preprocessing	14.6	2.3	22.3	65.2
	Mapping	49.2	8.7	75.5	
	Planning	1.4	1.6	2.2	
Reactive (32)	Preprocessing	13.8	1.3	72.3	19.1
	Planning	5.3	0.9	27.7	
Ours	Preprocessing	0.1	0.04	3.9	10.3 (2.6*)
	NN inference	10.1 (2.4*)	1.5 (0.8*)	93.0	
	Projection	0.08	0.01	3.1	
Ours (Onboard)	Preprocessing	0.2	0.1	0.4	41.6
	NN inference	38.9	4.5	93.6	
	Projection	2.5	1.5	6.0	

Our approach has substantially lower processing latency than both baselines; when network inference is performed on the GPU, our approach is 25.3 times faster than FastPlanner and 7.4 times faster than the Reactive baseline. When GPU inference is disabled, the network's latency increases by only 8 ms, and our approach is still much faster than both baselines. Moving from the desktop computer to the onboard embedded computing device, the network's forward pass requires 38.9 ms. Onboard, the total time to pass from the sensor reading to a plan is 41.6 ms, which corresponds to an update rate of about 24 Hz.

The effect of latency and sensor noise

We analyze the effect of sensor noise and planning latency in a controlled experiment. In this experiment, the quadrotor travels along a straight line at a constant forward speed and is required to laterally evade a single obstacle (a pole) while having only limited sensing range. This experimental setup was proposed by Falanga *et al.* (4) to understand the role of perception latency on the navigation ability of a robotic system subject to bounded inputs. Specifically, the authors derived an upper bound for the forward speed at which a robot can fly and avoid a single obstacle as a function of perception latency and sensing range. They modeled the robot as a point mass, which is a limited approximation for a quadrotor as it neglects the platform's rotational dynamics. We thus extend their formulation to account for the latency introduced by the rotational motion that is necessary to avoid the obstacle. A detailed description of the formulation can be found in the Supplementary Materials (section S5).

We set up the experiment by placing a quadrotor with an initial forward velocity v at a distance of 6 m from a pole with diameter 1.5 m. The quadrotor is modeled as a sphere with radius 0.2 m. According to our formulation, we compute a theoretical maximum speed—i.e., the speed at which the task is no longer feasible—for each method. The maximum speed depends on the sensing range, i.e., how far can an obstacle be accurately perceived; the latency of the visual sensor, i.e., the inverse of the frame rate; and the processing latency, i.e., the time to convert an observation into motor commands. Note that the latter varies between methods (c.f. the “Computational cost” section), and therefore, the theoretical maximum speed differs for each approach. For the details of the calculation, we refer the reader to section S5 in the Supplementary Materials. We then perform the controlled experiment with varying forward speeds v in the range of 3 to 13 m/s. We perform 10 experiments for each speed with all approaches and report the success rate. We run the experiment in two settings: (i) with ground-truth depth information to isolate the effect of latency on performance and (ii) with depth estimated by stereo matching (37) to analyze the effect of sensing errors to performance.

Ground-truth depth

Figure 5B1 illustrates the results of this experiment when perfect depth perception (Fig. 5A1) is available. All approaches can complete the task perfectly up to 5 m/s. However, even in these ideal conditions, the performance of the baselines drops for speeds beyond 5 m/s. This drop in performance for Reactive can be attributed to the fact that the finite library of motion primitives does not contain maneuvers that are aggressive enough to complete this task. Similarly, the performance degrades for FastPlanner as a result of suboptimal planning of actions. Although this baseline manages to map the obstacle in time, the planner frequently commands actions to stop the platform,

which leads to crashes when flying at high speeds. It is important to point out that the FastPlanner baseline was only demonstrated up to speeds of 3 m/s in the original work (18) and thus was not designed to operate at high speeds. Our approach can successfully avoid the obstacle without a single failure up to 7 m/s. For higher speeds, performance gracefully degrades to 60% at 10 m/s. This decrease in performance can be attributed to the sensitivity to imperfect network predictions when flying at high speed, where a single wrong action can lead to a crash.

Estimated depth


Whereas the previous experiments mainly focused on latency and the avoidance strategy, we now study the influence of imperfect sensory measurements on performance. We repeat the same experiment but provide all methods with depth maps that have been computed from the stereo pairs (Fig. 5A2). Figure 5B2 shows the results of this experiment. The baselines experience a substantial drop in performance compared with when provided with perfect sensory readings. FastPlanner completely fails for speeds of 5 m/s and beyond. This sharp drop in performance is due to the need for additional filtering of the noisy depth measurements that drastically increases the latency of the mapping stage. To reject outliers from the depth sensor, multiple observations (two/three) are required to completely map the obstacle. As a result, this baseline detects obstacles too late to be able to successfully evade them. Similarly, the performance of the Reactive baseline drops by 30% at 7 m/s. In contrast to the baselines, our approach is only marginally affected by the noisy depth readings with only a 10% drop in performance at 10 m/s, but no change in performance at lower speeds. This is because our policy, trained on depth from stereo, learns to account for common issues in the data such as discretization artifacts and missing values.

DISCUSSION

Existing autonomous flight systems are highly engineered and modular. The navigation task is usually split into sensing, mapping, planning, and control. The separation into multiple modules simplifies the implementation of engineered systems, enables parallel development of each component, and makes the overall system more interpretable. However, modularity comes at a high cost: The communication between modules introduces latency, errors compound across modules, and interactions between modules are not modeled. In addition, it is an open question whether certain subtasks, such as maintaining an explicit map of the environment, are even necessary for agile flight.

Our work replaces the traditional components of sensing, mapping, and planning with a single function that is represented by a neural network. This increases the system's robustness against sensor noise and reduces the processing latency. We demonstrate that our approach can reach speeds of up to 10 m/s in complex environments and reduces the failure rate at high speeds by up to 10 times when compared with the state of the art.

We achieve this by training a neural network to imitate an expert with privileged information in simulation. To cope with the complexity of the task and to enable seamless transfer from simulation to reality, we make several technical contributions. These include a sampling-based expert, a neural network architecture, and a training procedure, all of which take the task's multimodality into account. We also use an abstract, but sufficiently rich, input representation that considers real-world sensor noise. The combination of these innovations enables the training of robust navigation policies in



Ground Truth

Stereo

Real World

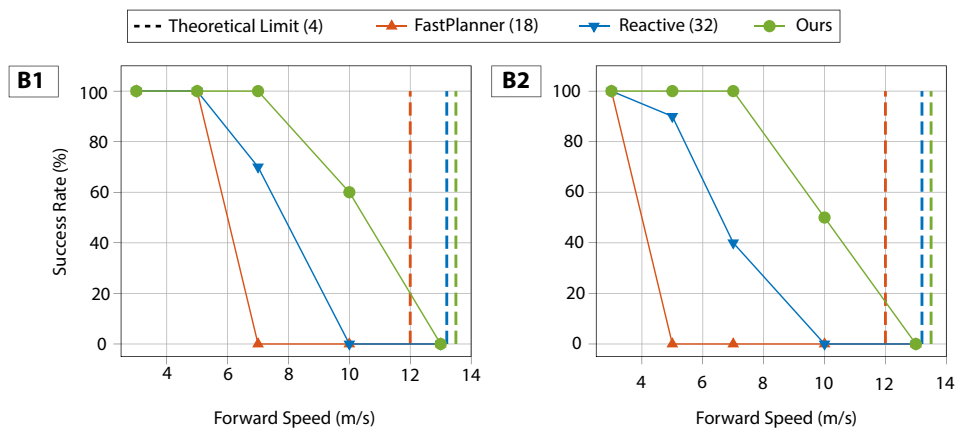


Fig. 5. The effect of sensor noise on performance. (A) When comparing RGB and depth images generated in simulation with images captured in the real world, we observe that the corresponding depth images are more similar than the RGB images. In addition, simulated depth estimated by stereo matching (A2) contains the typical failure cases of a depth sensor (A3), e.g., missing values and noise. (B) Results of the controlled experiment to study the effect of perception latency and noise on navigation ability. The experiment is performed on ground-truth depth (B1) and stereo depth (B2). Because each method has different processing latency (c.f. the “Computational cost” section), the theoretical maximum speed differs between approaches (dashed lines). Our method can fly closer to the theoretical limit than the baselines and is only minimally affected by the noise of stereo depth estimates.

simulation that can be directly transferred to diverse real-world environments without any fine-tuning on real data.

We see several opportunities for future work. Currently, the learned policy exhibits low success rates at average speeds of 10 m/s or higher in the real world. At these speeds, even our expert policy, despite having perfect knowledge of the environment, often fails to find collision-free trajectories. This is mainly due to the fact that, at speeds of 10 m/s or higher, feasible solutions require temporal consistency over a long time horizon and strong variations of the instantaneous flying speed as a function of the obstacle density. This requirement makes feasible trajectories extremely sparse in parameter space, resulting in intractable sampling. Engineering a more complex

expert to tackle this problem can be very challenging and might require specifically tailored heuristics to find approximate solutions. Therefore, we believe that this problem represents a big opportunity for the engineering requirements. The second reason for the performance drop at very high speeds is the mismatch between the simulated and physical drone in terms of dynamics and perception. The mismatch in dynamics is due to aerodynamics effects, motor delays, and dropping battery voltage. Therefore, we hypothesize that performance would benefit from increasing the fidelity of the simulated drone and making the policy robust to the unavoidable model mismatches. A third reason is perception latency: Faster sensors

can provide more information about the environment in a smaller amount of time and, therefore, can be used to provide more frequent updates, which improves the precision of the estimate. This could enable further reducing sensitivity to noise and promote a quicker understanding of the environment. Perception latency can possibly be reduced with event cameras (38), especially in the presence of dynamic obstacles (39).

Overall, our approach is a stepping stone toward the development of autonomous systems that can navigate at high speeds through previously unseen environments with only onboard sensing and computation. Combining our short-horizon controller for local obstacle avoidance with a long-term planner is a major opportunity for many robotics applications, including autonomous exploration, delivery, and cinematography.

MATERIALS AND METHODS

To perform agile flight through cluttered and previously unseen environments, we train a sensorimotor policy to predict receding-horizon trajectories from onboard sensor measurements and a reference trajectory. We assume that the reference trajectory is provided by a higher-level planning algorithm or the user. The reference is not necessarily collision free and only encodes the long-term goal of the platform. The agent is responsible to fly in the direction dictated by the reference while adapting its flight path to the environment.

An agent’s observation o consists of a depth image $\mathbf{d} \in \mathbb{R}^{640 \times 480}$, an estimate of the platform velocity $\mathbf{v} \in \mathbb{R}^3$ and attitude (expressed as a rotation matrix) $\mathbf{q} \in \mathbb{R}^9$, and a desired flight direction $\boldsymbol{\omega} \in \mathbb{R}^3$. The policy output is a set of motion hypotheses that are represented as receding-horizon trajectories with the corresponding estimated risk of collision. A model predictive controller then tracks the trajectory with the lowest collision probability and input cost. The controller is trained using privileged learning (30). Specifically, the policy is trained on demonstrations provided by a sampling-based planner that has access to information that is not available to the sensorimotor student: the precise knowledge of the 3D structure of the environment and the exact platform state. Figure 6 shows an overview of our method.

We collect demonstrations and perform training entirely in simulation. This trivially facilitates access to perfect 3D and state data for the privileged expert, enables the synthesis of unlimited training samples for any desired trajectory, and does not put the physical platform in danger. We use the Flightmare simulator (33) with the RotorS (34) Gazebo plugin and Unity as a rendering engine (35). Both the training and the testing environments are built by adding obstacles to the uneven ground of an out-of-the-box Unity environment.

To enable zero-shot transfer to real environments, special care has to be taken to minimize the domain shift of the (visual) input modalities from simulation to reality. To this end, we use depth as an abstract input representation that shows only a negligible domain shift from simulation to the real world (cf. Fig. 5A). Specifically, the

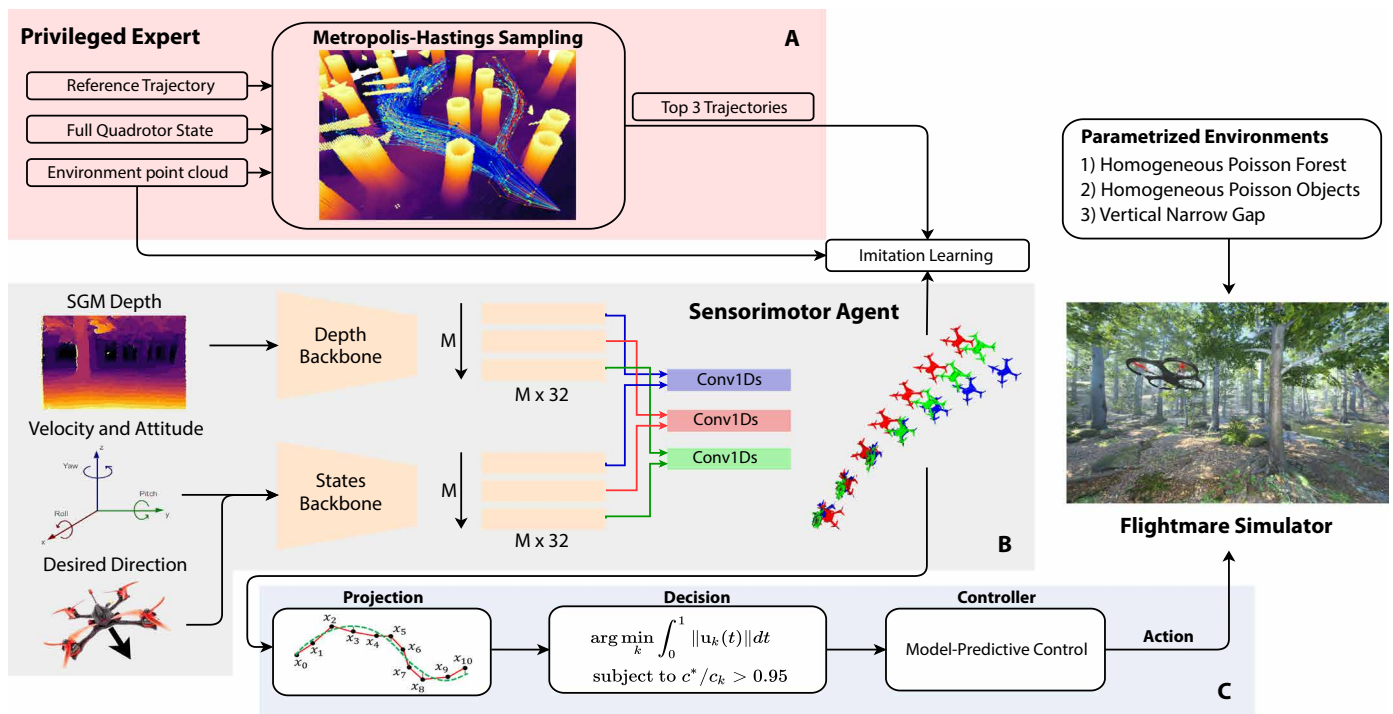


Fig. 6. Method overview. (A) Our offline planning algorithm computes a distribution of collision-free trajectories to follow a reference trajectory. The trajectories are computed with M-H sampling and are conditioned on complete 3D knowledge of the environment, which is represented by a point cloud. (B) A sensorimotor agent is trained with imitation learning to predict the best three trajectories from the estimated depth, the drone’s velocity and attitude, and the desired direction that encodes the goal. (C) The predictions are projected on the space of polynomial trajectories and ranked according to their predicted collision cost c_k . The trajectory with the lowest predicted cost c_k is then tracked with a model-predictive controller. If multiple trajectories have similar predicted cost (within a 5% range of the minimum $c^* = \min c_k$), the one with the smallest actuation cost is used.

sensorimotor agent is trained with depth images that have been computed using semiglobal matching (SGM) (37) from a simulated stereo camera pair. Because off-the-shelf depth sensors, such as the Intel RealSense 435 camera used on our physical platform, compute depth from stereo images using similar principles (40), this strategy ensures that the characteristics of the input align between simulation and the real world. As a result, the trained sensorimotor agent can be directly deployed in the real world. The next section presents both the privileged expert and the sensorimotor agent in detail.

The privileged expert

Our privileged expert is a sampling-based motion planning algorithm. The expert has perfect knowledge of the platform state and the environment (complete 3D map), both of which are only available in simulation. The expert generates a set of collision-free trajectories τ representing the desired state of the quadrotor $\mathbf{x}_{\text{des}} \in \mathbb{R}^{13}$ over the next second, starting from the current state of the drone, i.e., $\tau(0) = \mathbf{x}$. To do so, it samples from a probability distribution P that encodes distance from obstacles and proximity to the reference trajectory. Specifically, the distribution of collision-free trajectories $P(\tau | \tau_{\text{ref}}, C)$ is conditioned on the reference trajectory τ_{ref} and the structure of the environment in the form of a point cloud $C \in \mathbb{R}^{n \times 3}$. According to P , the probability of a trajectory τ is large if far from obstacles and close to the reference τ_{ref} . We define P as the following:

$$P(\tau | \tau_{\text{ref}}, C) = \frac{1}{Z} \exp(-c(\tau, \tau_{\text{ref}}, C)) \quad (1)$$

where $Z = \int_{\tau} P(\tau | \tau_{\text{ref}}, C)$ is the normalization factor and $c(\tau, \tau_{\text{ref}}, C) \in \mathbb{R}_+$ is a cost function indicating proximity to the reference and distance from obstacles. We define the trajectory cost function as

$$c(\tau, \tau_{\text{ref}}, C) = \int_0^1 \lambda_c C_{\text{collision}}(\tau(t)) + [\tau(t) - \tau_{\text{ref}}(t)]^T Q [\tau(t) - \tau_{\text{ref}}(t)] dt \quad (2)$$

where $\lambda_c = 1000$, Q is a positive semidefinite state cost matrix and $C_{\text{collision}}$ is a measure of the distance of the quadrotor to the points in C . We model the quadrotor as a sphere of radius $r_q = 0.2$ and define the collision cost as a truncated quadratic function of d_c , i.e., the distance between the quadrotor and the closest point in the environment

$$C_{\text{collision}}(\tau(t)) = \begin{cases} 0 & \text{if } d_c > 2r_q \\ -d_c^2 / r_q^2 + 4 & \text{otherwise} \end{cases} \quad (3)$$

The distribution P is complex because of the presence of arbitrary obstacles and frequently multimodal in cluttered environments because obstacles can be avoided in multiple ways. Therefore, the analytical computation of P is generally intractable.

To approximate the density P , the expert uses random sampling. We generate samples with the M-H algorithm (41) because it provides asymptotic convergence guarantees to the true distribution. To estimate P , the M-H algorithm requires a target score function $s(\tau) \propto P(\tau | \tau_{\text{ref}}, C)$. We define $s(\tau) = \exp(-c(\tau, \tau_{\text{ref}}, C))$, where $c(\cdot)$ is the cost of the trajectory τ . It is easy to show that this definition satisfies the conditions for the M-H algorithm to asymptotically estimate the target distribution P . Hence, the trajectories sampled with M-H will asymptotically cover all of the different modes of P . We point the interested reader to the Supplementary Materials

(section S6), for an overview of the M-H algorithm and its convergence criteria.

To decrease the dimension of the sampling space, we use a compact yet expressive representation of the trajectories τ . We represent τ as a cubic B-spline $\tau_{\text{bspline}} \in \mathbb{R}^{3 \times 3}$ curve with three control points and a uniform knot vector, enabling interpolation with high computational efficiency (42). Cubic B-splines are twice continuously differentiable and have a bounded derivative in a closed interval. Because of the differential flatness property of quadrotors (43), continuous and bounded acceleration directly translates to continuous attitude over the trajectory duration. This encourages dynamically feasible trajectories that can be tracked by a model-predictive controller accurately (43, 44). Therefore, instead of naively sampling the states of τ , we vary the shape of the trajectory by sampling the control points of the B-spline in a spherical coordinate system. In addition, for computational reasons, we discretize the trajectory at equally spaced time intervals of 0.1 s and evaluate the discrete version of Eq. 2. Specifically, we sample a total of 50 thousand trajectories using a Gaussian with a variance of 2, 5, 10 that increases every 16 thousand samples as the proposal distribution. Despite this efficient representation, the privileged expert cannot run in real time, given the large computational overhead introduced by sampling.

To bias the sampled trajectories toward obstacle-free regions, we replace the raw reference trajectory τ_{ref} in Eq. 2 with a global collision-free trajectory τ_{gbl} from start to goal that we compute using the approach of Liu *et al.* (13). As illustrated in fig. S6, conditioning sampling on τ_{gbl} practically increases the horizon of the expert and generates more conservative trajectories. An animation explaining our expert is available in Movie 1. After removing all generated trajectories in collision with obstacles, we select the three best trajectories with lower costs. Those trajectories are used to train the student policy.

The student policy

In contrast to the privileged expert, the student policy produces collision-free trajectories in real time with access only to onboard sensor measurements. These measurements include a depth image estimated with SGM (37), the platform velocity and attitude, and the desired direction of flight. The latter is represented as a normalized vector heading toward the reference point 1 s in the future with respect to the closest reference state. We hypothesize that this information is sufficient for generating the highest probability samples of the distribution $P(\tau | \tau_{\text{ref}}, C)$ without actually having access to the point cloud C . There are two main challenges to accomplishing this: (i) The environment is only partially observable from noisy sensor observations and (ii) the distribution P is in general multimodal. Multiple motion hypotheses with high probabilities can be available, but their average can have a very low probability.

We represent the policy as a neural network that is designed to be able to cope with these issues. The network consists of an architecture with two branches that produce a latent encoding of visual, inertial, and reference information and outputs $M = 3$ trajectories and their respective collision cost. We use a pretrained MobileNet-V3 architecture (45) to efficiently extract features from the depth image. The features are then processed by a 1D convolution to generate M feature vectors of size 32. The current platform's velocity and attitude are then concatenated with the desired reference direction and processed by a four-layer perceptron with [64, 32, 32, 32] hidden nodes and LeakyReLU activations. We again use 1D convolutions to create a 32-dimensional feature vector for each mode. The visual

and state features are then concatenated and processed independently for each mode by another four-layer perceptron with [64, 128, 128] hidden nodes and LeakyReLU activations. The latter predicts, for each mode, a trajectory τ and its collision cost. In summary, our architecture receives as input a depth image $\mathbf{d} \in \mathbb{R}^{640 \times 480}$, the platform's velocity $\mathbf{v} \in \mathbb{R}^3$, the drone's attitude expressed as a rotation matrix $\mathbf{q} \in \mathbb{R}^9$, and reference direction $\omega \in \mathbb{R}^3$. From this input, it predicts a set \mathcal{T}_n of trajectories and their relative collision cost, i.e., $\mathcal{T}_n = \{(\tau_n^k, c_k) \mid k \in [0, 1, \dots, M-1]\}$, where $c_k \in \mathbb{R}_+$. Differently from the privileged expert, the trajectory predicted by the network does not describe the full state evolution but only its position component, i.e., $\tau_n^k \in \mathbb{R}^{10 \times 3}$. Specifically, the network trajectory τ_n^k is described by

$$\tau_n^k = [\mathbf{p}(t_i)]_{i=1}^{10}, t_i = \frac{i}{10} \quad (4)$$

where $\mathbf{p}(t_i) \in \mathbb{R}^3$ is the drone's position at time $t = t_i$ relative to its current state \mathbf{x} . This representation is more general than the B-spline with three control points used by the sampling-based planner, and it was preferred to the latter representation to avoid the computational costs of interpolation at test time. We train the neural network with supervised learning on the three trajectories with lowest cost found by the expert. To account for the multihypotheses prediction, we minimize the following relaxed winner-takes-all (R-WTA) loss for each sample

$$\text{R-WTA}(\mathcal{T}_e, \mathcal{T}_n) = \sum_{i=0}^{|\mathcal{T}_e|} \sum_{k=0}^{|\mathcal{T}_n|} \alpha(\tau_{e,p}^i, \tau_n^k) \|\tau_{e,p}^i - \tau_n^k\|^2 \quad (5)$$

where \mathcal{T}_e and \mathcal{T}_n are the set of expert and network trajectories, $\tau_{e,p}$ denotes the position component of τ_e , and $\alpha(\cdot)$ is defined as

$$\alpha(\tau_{e,p}^i, \tau_n^k) = \begin{cases} 1 - \epsilon & \text{if } \|\tau_{e,p}^i - \tau_n^k\|^2 \leq \|\tau_{e,p}^j - \tau_n^k\|^2 \forall j \neq i \\ \frac{\epsilon}{M-1} & \text{otherwise} \end{cases} \quad (6)$$

A necessary condition for Eq. 5 to be minimized is that the network predictions \mathcal{T}_n correspond to the centroids of a Voronoi tessellation of expert labels \mathcal{T}_e (46). However, Eq. 5 is not differentiable because it performs a hard assignment between predictions and labels. Therefore, we relax the assignment with a small value ϵ (Eq. 6) to optimize Eq. 5 with gradient descent. Intuitively, an expert trajectory $\tau_e^i \in \mathcal{T}_e$ is associated with the closest network trajectory $\tau_n^k \in \mathcal{T}_n$ with a weight of $1 - \epsilon = 0.95$ and with $\epsilon/(M-1) = 0.025$ to all remaining hypotheses. This formulation, proposed by Rupprecht *et al.* (46), prevents mode collapse and was shown to outperform other popular approaches for multimodal learning such as mixture density networks (47). In addition, the predicted collision cost c_k is trained with supervised learning on the ground-truth cost $C_{\text{collision}}(\tau_n^k)$ computed according to Eq. 3. In summary, the final training loss for each sample is equal to

$$\mathcal{L}(\mathcal{T}_e, \mathcal{T}_n) = \lambda_1 \text{R-WTA}(\mathcal{T}_e, \mathcal{T}_n) + \lambda_2 \sum_{k=0}^{|\mathcal{T}_n|} (c_k - C_{\text{collision}}(\tau_n^k))^2 \quad (7)$$

where $\lambda_1 = 10$ and $\lambda_2 = 0.1$ were empirically found to equalize the magnitudes of the two terms. This loss is averaged over a mini batch

of eight samples and minimized with the Adam optimizer (48) and a learning rate of 1×10^{-3} .

At test time, we retrieve the full state information from the predicted trajectories by projecting them on the space of order-5 polynomials for each axis independently. This representation enforces continuity in position, velocity, and acceleration with respect to the current state and facilitates dynamic feasibility due to differential flatness (43). Considering for example the x axis, we define the polynomial projection $\mu_x(t) = \mathbf{a}_x^\top \cdot \mathbf{T}(t)$, where $\mathbf{a}_x^\top = [a_0, a_1, \dots, a_5]$ and $\mathbf{T}(t)^\top = [1, t, \dots, t^5]$. The projection term \mathbf{a}_x is found by solving the following optimization problem

$$\begin{aligned} & \underset{\mathbf{a}_x}{\text{minimize}} \sum_{i=1}^{10} \left(\tau_{n,x}^{k,i} - \mathbf{a}_x^\top \cdot \mathbf{T}\left(\frac{i}{10}\right) \right)^2 \\ & \text{subject to } s_x(0) - \mathbf{a}_x^\top \cdot \mathbf{T}(0) = 0 \\ & \dot{s}_x(0) - \mathbf{a}_x^\top \cdot \dot{\mathbf{T}}(0) = 0 \\ & \ddot{s}_x(0) - \mathbf{a}_x^\top \cdot \ddot{\mathbf{T}}(0) = 0 \end{aligned} \quad (8)$$

where $\tau_{n,x}^{k,i}$ is the x component of the i th element of τ_n^k and $s_x(0)$, $\dot{s}_x(0)$, $\ddot{s}_x(0)$ are the x position of the quadrotor and its derivatives obtained from the current state estimate. The latter corresponds to the ground-truth state when deployed in simulation and to the state estimate computed by the Intel RealSense T265 when deployed on the physical platform. To reduce abrupt changes in velocity during flight, which would lead to strong pitching motion, we additionally constrain the polynomial's average speed to a desired value v_{des} . To do so, we scale the time t of polynomial $\mu_x(t)$ by a factor $\beta = v_{\text{des}}/v_{\mu}^x$, i.e., $t' = \beta t$, where $v_{\mu}^x = |\mu(1) - \mu(0)|$.

Once all predicted trajectories are projected, we select one for execution. To do so, we select trajectories with $c^*/c_k \geq 0.95$ ($c^* = \min c_k$) and compute their input costs according to Mellinger and Kumar (43). The one with the lowest input costs is tracked by a model-predictive controller (44). Intuitively, this choice enforces temporal continuity in the trajectories. For example, when dodging an obstacle to the right, we do not want to steer toward the left at the next iteration, unless strictly necessary because of the appearance of another obstacle.

Training environments

We build custom environments in the Flightmare simulator (33) to collect training data. All environments are built by spawning items on the uneven empty ground of an off-the-shelf Unity environment. We spawn items belonging to two categories: simulated trees, available off-the-shelf, and a set of convex shapes such as ellipsoids, cuboids, and cylinders. Sample observations from these environments are available in Fig. 7. Additional training samples, together with observations collected in the simulated testing environments, are available in fig. S4. The dimensions of these shapes are randomized according to a continuous uniform random distribution with $x \in \mathcal{U}(0.5, 4)$, $y \in \mathcal{U}(0.5, 4)$, and $z \in \mathcal{U}(0.5, 8)$. Training environments are created by spawning either of the two categories of items according to a homogeneous Poisson point process with intensity δ .

We generate a total of 850 environments by uniform randomization of the following two quantities: item category, i.e., trees or shapes, and the intensity $\delta \in \mathcal{U}(4, 7)$ with $\delta \in \mathbb{N}_+$. For each environment,



Fig. 7. Method validation. Without initialization of the sampler on a global trajectory (LocalPlan), multimodal training (UniModal), or training on SGM depth image (RGB), performance drops. Predicting a single point in the future instead of a full trajectory (SinglePoint) has similar effects on performance. Ours consistently outperforms all the ablated versions of the system. (B and C) Training environments. We build training environments by spawning trees or convex shapes according to a homogeneous Poisson point process with varying intensities. Additional samples of the training and testing simulation environments are available in fig. S4.

we compute a global collision-free trajectory τ_{gbl} from the starting location to a point 40 m in front of it. The trajectory τ_{gbl} is not observed by the student policy, but only by the expert. The student is only provided with a straight, potentially not collision-free, trajectory from start to end to convey the goal.

To assure sufficient coverage of the state space, we use the dataset aggregation strategy (DAGger) (49). This process consists of rolling out the student policy and labeling the visited states with the expert policy. To avoid divergence from τ_{gbl} and prevent crashes in the early stages of training, we track a trajectory predicted by the student only if the drone's distance from the closest point on τ_{gbl} is smaller than a threshold ξ , initialized to zero. Otherwise, we directly track τ_{gbl} with a model-predictive controller. Every 30 environments, the student is retrained on all available data and the threshold ξ set to $\xi' = \min(\xi + 0.25, 6)$. Aggregating data over all environments results in a dataset of about 90 thousand samples. For the narrow-gap experiments, we fine-tune the student policy on 100 environments created by adding a 50-m-long wall with a single vertical gap of random width $w_g \in \mathcal{U}(0.7, 1.2)$ m in the center. In those environments, the drone starts at a 10-m distance from the wall with a randomized lateral offset $l \in \mathcal{U}(-5, 5)$ from the gap. We collect about 10 thousand training samples from those environments.

We evaluate the trained policies in simulation on environments not seen during training but coming from the same distributions. The same policies are then used to control the physical platforms in real-world environments. When deployed in simulation, we estimate depth with SGM (37) from a simulated stereo pair and use the ground-truth state of the platform. Conversely, on the physical platform, depth is estimated by an off-the-shelf Intel RealSense 435, and state estimation is performed by an Intel RealSense T265. More details on our experimental platform are available in section S3.

Method validation

Our approach is based on several design choices that we validate in an ablation study. We ablate the following components: (i) the use of global planning to initialize the sampling of the privileged expert, (ii) the use of depth as an intermediate representation for action,

and (iii) multimodal network prediction. The results in Fig. 7 show that all components are important and that some choices have a larger effect than others. Our study indicates that depth perception plays a fundamental rule for high-speed obstacle avoidance: When training on color images (RGB), performance drops substantially. This is in line with previous findings (28, 50) showing that intermediate image representations have a strong positive effect on performance when sufficiently informative for the task. Not accounting for multimodal trajectory prediction (UniModal) is also detrimental for performance. This is because the loss pushes predicted trajectories toward the average of the expert trajectories, which is often in collision with obstacles. Not initializing the sampling of expert trajectories on a global plan (LocalPlan) is not important for low-speed flight but plays an important role for success at higher speeds. Biasing the motion away from regions that are densely populated by obstacles is particularly important for high-speed flight where little slack is available for reacting to unexpected obstacles. In addition, we compare our output representation, i.e., a trajectory 1 s in the future, to the less complex output representation in prior work (51), i.e., a single waypoint in the future (SinglePoint). The results indicate that the limited representation power of the latter representation causes performance to drop, especially at high speeds.

SUPPLEMENTARY MATERIALS

www.science.org/doi/10.1126/scirobotics.abg5810
 Sections S1 to S6
 Figs. S1 to S6
 Tables S1 to S3
 Movies S1 and S2
 References (S2–S5)

REFERENCES AND NOTES

1. J. Verbeke, J. D. Schutter, Experimental maneuverability and agility quantification for rotary unmanned aerial vehicle. *Int. J. Micro Air Veh.* **10**, 3–11 (2018).
2. E. Ackerman, Ai-powered drone learns extreme acrobatics (IEEE Spectrum, 2020); <https://spectrum.ieee.org/automaton/robotics/drones/ai-powered-drone-extreme-acrobatics>.
3. J. Hilf, K. Umbach, The commercial use of drones. *Comput. Law Rev. Int.* **16**, 3 (2015).

4. D. Falanga, S. Kim, D. Scaramuzza, How fast is too fast? The role of perception latency in high-speed sense and avoid. *IEEE Robot. Autom. Lett.* **4**, 1884–1891 (2019).
5. L. Heng, D. Honegger, G. H. Lee, L. Meier, P. Tanskanen, F. Fraundorfer, M. Pollefeys, Autonomous visual mapping and exploration with a micro aerial vehicle. *J. Field Robot.* **31**, 654–675 (2014).
6. S. Saedi, C. Thibault, M. Trentini, H. Li, 3D mapping for autonomous quadrotor aircraft. *Unmanned Syst.* **5**, 181–196 (2017).
7. D. Scaramuzza, M. Achtelik, L. Doitsidis, F. Fraundorfer, E. B. Kosmatopoulos, A. Martinelli, M. W. Achtelik, M. Chli, S. A. Chatzichristofis, L. Kneip, D. Gurdan, L. Heng, G. H. Lee, S. Lynen, M. Pollefeys, A. Renzaglia, R. Siegwart, J. C. Stumpf, P. Tanskanen, C. Troiani, S. Weiss, L. Meier, Vision-controlled micro flying robots: From system design to autonomous navigation and mapping in GPS-denied environments. *IEEE Robotics Autom. Mag.* **21**, 26–40 (2014).
8. M. Faessler, F. Fontana, C. Forster, E. Mueggler, M. Pizzoli, D. Scaramuzza, Autonomous, vision-based flight and live dense 3D mapping with a quadrotor micro aerial vehicle. *J. Field Robotics* **33**, 431–450 (2016).
9. M. Blösch, S. Weiss, D. Scaramuzza, R. Siegwart, Vision based MAV navigation in unknown and unstructured environments, in *2010 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2010), pp. 21–28.
10. A. Bry, N. Roy, Rapidly-exploring random belief trees for motion planning under uncertainty, in *2011 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2011), pp. 723–730.
11. C. Richter, A. Bry, N. Roy, Polynomial trajectory planning for aggressive quadrotor flight in dense indoor environments, in *Proceedings of the International Symposium on Robotics Research (ISRR)* (ISRR, 2013), pp. 649–666.
12. R. Allen, M. Pavone, A real-time framework for kinodynamic planning with application to quadrotor obstacle avoidance, in *AIAA Guidance, Navigation, and Control Conference* (AIAA, 2016), p. 1374.
13. S. Liu, K. Mohta, N. Atanasov, V. Kumar, Search-based motion planning for aggressive flight in SE(3). *IEEE Robot. Autom. Lett.* **3**, 2439–2446 (2018).
14. H. Oleynikova, Z. Taylor, M. Fehr, R. Siegwart, J. I. Nieto, Voxblox: Incremental 3d euclidean signed distance fields for on-board MAV planning, in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2017), pp. 1366–1373.
15. T. Ozaslan, G. Loianno, J. Keller, C. J. Taylor, V. Kumar, J. M. Wozencraft, T. Hood, Autonomous navigation and mapping for inspection of penstocks and tunnels with mavs. *IEEE Robot. Autom. Lett.* **2**, 1740–1747 (2017).
16. K. Mohta, M. Watterson, Y. Mulgaonkar, S. Liu, C. Qu, A. Makineni, K. Saultner, K. Sun, A. Zhu, J. Delmerico, K. Karydis, N. Atanasov, G. Loianno, D. Scaramuzza, K. Daniilidis, C. J. Taylor, V. Kumar, Fast, autonomous flight in gps-denied and cluttered environments. *J. Field Robot.* **35**, 101–120 (2018).
17. A. J. Barry, P. R. Florence, R. Tedrake, High-speed autonomous obstacle avoidance with pushbroom stereo. *J. Field Robot.* **35**, 52–68 (2018).
18. B. Zhou, F. Gao, L. Wang, C. Liu, S. Shen, Robust and efficient quadrotor trajectory generation for fast autonomous flight. *IEEE Robot. Autom. Lett.* **4**, 3529–3536 (2019).
19. M. Ryll, J. Ware, J. Carter, N. Roy, Efficient trajectory planning for high speed flight in unknown environments, in *2019 International Conference on Robotics and Automation (ICRA)* (IEEE, 2019), pp. 732–738.
20. J. Tordesillas, B. T. Lopez, J. P. How, FASTER: Fast and safe trajectory planner for flights in unknown environments, in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2019), pp. 1934–1940.
21. T. Cieslewski, E. Kaufmann, D. Scaramuzza, Rapid exploration with multi-rotors: A frontier selection method for high speed flight, in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2017), pp. 2135–2142.
22. Z. Zhang, D. Scaramuzza, Perception-aware receding horizon navigation for MAVs, in *IEEE International Conference on Robotics and Automation (ICRA)* (IEEE 2018), pp. 2534–2541.
23. G. Loianno, D. Scaramuzza, Special issue on future challenges and opportunities in vision-based drone navigation. *J. Field Robot.* **37**, 495–496 (2020).
24. S. Ross, N. Melik-Barkhudarov, K. S. Shankar, A. Wendel, D. Dey, J. A. Bagnell, M. Hebert, Learning monocular reactive UAV control in reduced natural environments, in *2013 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2013), pp. 1765–1772.
25. F. Sadeghi, S. Levine, CAD2RL: Real single-image flight without a single real image, in *Proceedings of the Robotics: Science and Systems (RSS) Conference*, N. M. Amato, S. S. Srinivasa, N. Ayanian, S. Kuindersma, Eds. (2017).
26. D. Gandhi, L. Pinto, A. Gupta, Learning to fly by crashing, in *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2017), pp. 3948–3955.
27. A. Loquercio, A. I. Maqueda, C. R. del-Blanco, D. Scaramuzza, Dronet: Learning to fly by driving. *IEEE Robotics Autom. Lett.* **3**, 1088–1095 (2018).
28. E. Kaufmann, A. Loquercio, R. Ranftl, M. Müller, V. Koltun, D. Scaramuzza, Deep drone acrobatics, in *Proceedings of the Robotics: Science and Systems (RSS) Conference*, M. Toussaint, A. Bicchi, T. Hermans, Eds. (2020).
29. M. Müller, A. Dosovitskiy, B. Ghanem, V. Koltun, Driving policy transfer via modularity and abstraction, in *Proceedings of the 2nd Conference on Robot Learning* (PMLR, 2018), pp. 1–15.
30. D. Chen, B. Zhou, V. Koltun, P. Krährenbühl, Learning by cheating,” in *Proceedings of the Conference on Robot Learning* (PMLR, 2019), pp. 66–75.
31. B. Zhou, J. Pan, F. Gao, S. Shen, RAPTOR: Robust and perception-aware trajectory replanning for quadrotor fast flight, *CoRR*. arXiv:2007.03465 [cs.RO] (6 July 2020).
32. P. Florence, J. Carter, R. Tedrake, Integrated perception and control at high speed: Evaluating collision avoidance maneuvers without maps, in *Algorithmic Foundations of Robotics XII*, K. Goldberg, P. Abbeel, K. Bekris, L. Miller, Eds. (Springer International Publishing, 2020), vol. 13, pp. 304–319.
33. Y. Song, S. Naji, E. Kaufmann, A. Loquercio, D. Scaramuzza, Flightmare: A flexible quadrotor simulator, in *4th Conference on Robot Learning* (PMLR, 2020).
34. F. Furrer, M. Burri, M. Achtelik, R. Siegwart, Rotors—A modular gazebo mav simulator framework, in *Robot Operating System (ROS)*, A. Koubaa, Ed. (Springer, 2016), pp. 595–625.
35. A. Juliani, V.-P. Berges, E. Teng, A. Cohen, J. Harper, C. Elion, C. Goy, Y. Gao, H. Henry, M. Mattar, D. Lange, Unity: A general platform for intelligent agents, arXiv:1809.02627 [cs.LG] (7 September 2018).
36. S. Karaman, E. Frazzoli, High-speed flight in an ergodic forest, in *2012 IEEE International Conference on Robotics and Automation* (IEEE, 2012), pp. 2899–2906.
37. H. Hirschmuller, Stereo processing by semiglobal matching and mutual information. *IEEE Trans. Pattern Anal. Mach. Intell.* **30**, 328–341 (2007).
38. G. Gallego, T. Delbruck, G. Orchard, C. Bartolozzi, B. Tabla, A. Censi, S. Leutenegger, A. Davison, J. Conradt, K. Daniilidis, D. Scaramuzza, Event-based vision: A survey, *IEEE Trans. Pattern Anal. Machine Intell.* (2020); available at <https://arxiv.org/abs/1904.08405>.
39. D. Falanga, K. Kleber, D. Scaramuzza, Dynamic obstacle avoidance for quadrotors with event cameras. *Sci. Robot.* **5**, aaz9712 (2020).
40. L. Keselman, J. Iselin Woodfill, A. Grunnet-Jepsen, A. Bhowmik, Intel realsense stereoscopic depth cameras, in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshops* (IEEE, 2017), pp. 1–10.
41. W. K. Hastings, Monte carlo sampling methods using markov chains and their applications. *Biometrika* **57**, 97–109 (1970).
42. J. Gallier, *Curves and Surfaces in Geometric Modeling: Theory and Algorithms* (Morgan Kaufmann Publishers Inc., 1999).
43. D. Mellinger, V. Kumar, Minimum snap trajectory generation and control for quadrotors, in *2011 IEEE International Conference on Robotics and Automation (ICRA)* (IEEE, 2011), pp. 2520–2525.
44. D. Falanga, P. Fehrer, P. Lu, D. Scaramuzza, Pampc: Perception-aware model predictive control for quadrotors, in *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE, 2018), pp. 1–8.
45. A. Howard, M. Sandler, B. Chen, W. J. Wang, L. C. Chen, M. X. Tan, G. Chu, V. Vasudevan, Y. K. Zhu, R. M. Pang, H. Adam, Q. Le, Searching for mobileNetV3, in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (IEEE, 2019), pp. 1314–1324.
46. C. Rupprecht, I. Laina, R. S. DiPietro, M. Baust, Learning in an uncertain world: Representing ambiguity through multiple hypotheses, in *Proceedings of the 2017 IEEE International Conference on Computer Vision (ICCV)* (IEEE, 2017), pp. 3611–3620.
47. C. M. Bishop, *Mixture Density Networks* (Aston Univ., 1994).
48. D. P. Kingma, L. J. Ba, Adam: A method for stochastic optimization, in *Proceedings of the International Conference on Learning Representations (ICLR)* (ICLR, 2015).
49. S. Ross, G. Gordon, D. Bagnell, A reduction of imitation learning and structured prediction to no-regret online learning, in *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics* (PMLR, 2011), pp. 627–635.
50. B. Zhou, P. Krährenbühl, V. Koltun, Does computer vision matter for action? *Sci. Robot.* **4**, aaw6661 (2019).
51. A. Loquercio, E. Kaufmann, R. Ranftl, A. Dosovitskiy, V. Koltun, D. Scaramuzza, Deep drone racing: From simulation to reality with domain randomization. *IEEE Trans. Robot.* **36**, 1–14 (2019).
52. E. Tomppa, *Models and methods for analysing spatial patterns of trees*, thesis, Finnish Forest Research Institute (1986).
53. D. Hernandez-Juarez, A. Chacón, A. Espinosa, D. Vázquez, J. C. Moure, A. M. López, Embedded real-time stereo estimation via semi-global matching on the GPU, in *International Conference on Computational Science 2016, ICCS 2016*, 6 to 8 June 2016, San Diego, California, USA, 2016.
54. C. Andrieu, N. de Freitas, A. Doucet, M. I. Jordan, An introduction to MCMC for machine learning. *Mach. Learn.* **50**, 5–43 (2003).
55. G. O. Roberts, A. F. M. Smith, Simple conditions for the convergence of the Gibbs sampler and Metropolis-Hastings algorithms. *Stoch. Process. Their Appl.* **49**, 207–216 (1994).

Acknowledgments: We thank T. Längle, M. Sutter, Y. Song, C. Pfeiffer, L. Bauersfeld, and N. Aepli for contributions to the drone design, the simulation environment, the support during real-world experiments, and preparation of multimedia material. **Funding:** This work was

supported by the Intel Network on Intelligent Systems, Armasuisse, the National Centre of Competence in Research (NCCR) Robotics through the Swiss National Science Foundation (SNSF), and by the European Research Council (ERC) under Grant Agreement 864042 (AGILEFLIGHT). **Author contributions:** A.L. and E.K. formulated the main ideas, implemented the system, performed all experiments, and wrote the paper. R.R. and M.M. contributed to the experimental design, data analysis, and paper writing. V.K. and D.S. provided funding and contributed to the design and analysis of experiments. **Competing interests:** The authors declare that they have no competing interests. **Data and materials availability:** All data needed to evaluate the conclusions in the paper are present in the paper or the Supplementary Materials. Videos of the experiments and the codebase of the work can be

found at the project webpage <http://rpg.ifi.uzh.ch/AgileAutonomy.html> and <https://doi.org/10.5281/zenodo.5517791>, respectively.

Submitted 15 February 2021
Accepted 13 September 2021
Published 6 October 2021
10.1126/scirobotics.abg5810

Citation: A. Loquercio, E. Kaufmann, R. Ranftl, M. Müller, V. Koltun, D. Scaramuzza, Learning high-speed flight in the wild. *Sci. Robot.* **6**, eabg5810 (2021).

Learning high-speed flight in the wild

Antonio Loquercio, Elia Kaufmann, René Ranftl, Matthias Müller, Vladlen Koltun, and Davide Scaramuzza

Sci. Robot. **6** (59), eabg5810. DOI: 10.1126/scirobotics.abg5810

View the article online

<https://www.science.org/doi/10.1126/scirobotics.abg5810>

Permissions

<https://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of service](#)

Science Robotics (ISSN 2470-9476) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. The title *Science Robotics* is a registered trademark of AAAS.

Copyright © 2021 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works