MIT Mobility Initiative Final Report

Enabling Equitable Access to Advanced Air Mobility Services: Routing with preferences for uncrewed aerial vehicles

Geoffrey Ding and Hamsa Balakrishnan

August 2026

Contents

Executive Summary **2**

1 Introduction **2**

2 Methodology **3**
2.1 Modeling preferences . . . 3
2.2 Properties of preferred segments . . . 3
2.3 The price of preference . . . 4
2.4 Obstacles . . . 5

3 Example in a synthetic environment **6**
3.1 Setup . . . 6
3.2 Results . . . 7

4 Drone as first responder (DFR) case study **7**
4.1 Setup . . . 9
4.2 Routing . . . 9
4.3 Results . . . 11
4.4 Travel time comparison . . . 12

5 Limitations and stakeholder implications **14**
5.1 Limitations . . . 14
5.2 Implications for operators . . . 14
5.3 Implications for policymakers . . . 15

6 Next steps **15**

Executive summary

As the number of uncrewed aerial vehicles (UAVs), or drones, flights increase, there has been growing public concern about safety, privacy, and noise. These risks may be reduced by routing along paths with less human activity such as waterways and railways, which we term *preferred segments*.

By varying the number of preferred segments and how strongly they are preferred, we characterize how preferences affect routes and quantify how they impact operational efficiency. When there are only a few preferred segments in a synthetic environment, flights may need to take significant detours to minimize cost, especially when preferences are strong. However, this drawback diminishes as the number of preferred segments increases.

We also provide an illustrative drone as first responder (DFR) use case, finding find that impacts are not as bad in practice as they could be in theory. Deploying representative aircraft from just 5 out of 39 facilities has a median reduction in response times of 18 to 30 percent compared to ground vehicles, depending on the degree of preferences. However, the benefit does vary by incident, with between 26 and 41 percent of responses *slower* than if a ground vehicle had been deployed. Therefore, dispatchers must be aware of how different situations benefit differently from DFR.

For operators, the trade-off between routing with preferences and taking the shortest path depends on the urgency of their task. Emergency response will likely aim to minimize travel time, while drone deliveries may be more amenable to detours to mitigate safety, privacy, and noise risks. On the other hand, regulators must be cautious that their policies are not inadvertently counterproductive by concentrating flights outside of preferred regions.

1 Introduction

The use of uncrewed aerial vehicles (UAVs), or drones, for recreation, commerce, and public safety has been rapidly increasing. Beyond the technical challenges of drone operations at scale, we must also address societal concerns including, but not limited to, safety, privacy, and noise.

Safety is paramount to public acceptance of widespread drone use, with ground impacts and midair collisions being leading safety risks [15]. As a result, regulators around the world today prioritize safety [10], such as the Federal Aviation Administration (FAA) restricting on operations over people and moving vehicles [6]. Despite this, prior work has found privacy to be more important to the public than safety [5, 1], especially the risk of unwanted surveillance from the cameras that many drones carry [7, 14, 4]. While noise pollution is already an issue in many urban areas, a survey of drone noise emissions and their effects on people finds that "drone noise is substantially more annoying than road traffic or aircraft noise at the same level" due to its acoustic characteristics [12]. As a result, noise impacts have been considered in routing [9, 8] and facility location [13] decisions.

Based on the population distribution and travel patterns in a region, drone operations along certain routes may pose lower safety, privacy, and noise risks. We study the implications of path planning for drone operations when there are these kinds of preferences. Because it is most feasible to respect these preferences when operators have flexibility over their routing decisions, we focus on drone applications that require coverage of a wide area from a small number of bases, such as package delivery and drone as first responder (DFR). This is in contrast to functions that are necessarily constrained to specific locations, such as aerial photography or infrastructure inspections.

The remainder of this work is organized as follows. In Sec. 2, we introduce routing with preferences, including with obstacles, and offer some theoretical results. Then, we study a synthetic environment in Sec. 3 and a DFR application using real data from the City of Detroit in Sec. 4. Section 5 discusses the implications of our results for drone operators and regulators as well as the limitations of our approach. We offer some next steps in Sec. 6.

2 Methodology

Suppose drone flights are preferable to a general baseline in some regions, such as along waterways or roadways. These areas offer lower risks to public safety and privacy, and may be less disruptive, than flying over more populated areas.

2.1 Modeling preferences

We focus on linear preferred regions and represent them with line segments that can be traversed at a lower cost per unit distance. Each *preferred segment*, or simply segment, has a *discount factor* $\alpha \in [0, 1]$, representing the relative cost of travel along the segment compared to the baseline.

The discount factor can be viewed as a measure of urgency. Values closer to 1 are appropriate for time-sensitive functions like DFR, which requires vehicles to reach their destinations as quickly as possible. By contrast, values closer to 0 may be acceptable for applications like package delivery, which have milder consequences for detours or delays.

Alternatively, α can be viewed as the inverse of a inverse speedup factor. If the baseline travel speed is normalized to 1, a segment can be traversed at speed $1/\alpha \geq 1$. Minimizing travel time in this scheme is equivalent to minimizing cost in the previous one.

We call the path with lowest cost, which considers preferences, the *optimal* path. The path of least length, which ignores preferences, is the *shortest* path.

2.2 Properties of preferred segments

A key property of a segment is its *critical incident angle*, or critical angle, given by $\theta^* = \sin^{-1} \alpha$ and measured relative to the direction perpendicular to the segment [11]. Figure 1 shows optimal paths from the center of the region to randomly distributed points. The black line represents a preferred segment, and blue lines represent paths that do not intersect the segment. Several properties of optimal paths relative to the critical angle can be seen:

1. An optimal path that intersects the segment does not use it if and only if the angle of incidence is less than θ^*, as illustrated by the orange paths.
2. An optimal path that enters or exits the interior of the segment must do so at exactly the critical incident angle, as illustrated by the green paths.
3. An optimal path that enters or exits an endpoint of the segment must do so at angle greater than or equal to the critical incident angle away from the segment, as illustrated by the red paths.

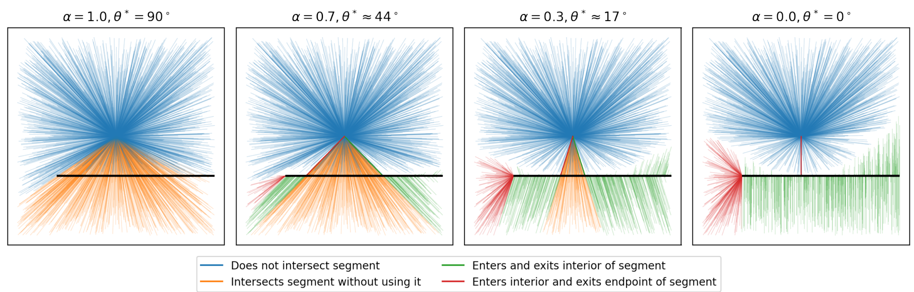

Figure 1: Optimal paths with a single segment in black at various critical incidence angles θ^*, defined relative to the segment normal. Each colored line represents one optimal path from origin to destination.

When $\alpha = 1$, there is no discount, the segment will never be used in any optimal path, and the shortest path is optimal. Traveling along a segment with $\alpha = 0$ incurs no cost, so any optimal path that crosses the segment will use it.

The properties of preferred segments allow us to construct a graph that contains all possible optimal paths, or the *critical graph* [11]. Then, solving a least-cost path problem returns the optimal path. In Fig. 2, we use the critical graph to find optimal paths and compare them to shortest paths across several values of α in a synthetic environment.

2.3 The price of preference

Optimal paths using preferred segments may be desirable for several reason, but deviating from the shortest path incurs additional travel time, energy consumption, or other costs. Thus, we define the *price of preference* for a given origin, destination, and environment as the ratio between the length of the optimal path and that of the shortest path.

The price of preference depends on the geometry of the segments. If the minimum discount factor across all segments is $\alpha_{\min}$, then the price of preference

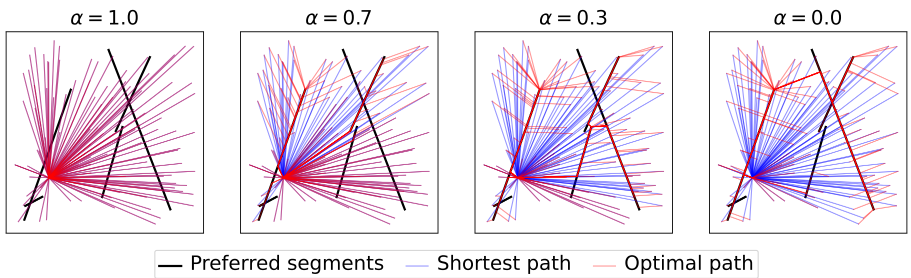

Figure 2: Optimal and shortest paths with various α.

is at most $1/\alpha_{\min}$, but this is almost always an overestimate. For example, suppose there are three segments forming an equilateral triangle, each with a discount of $\alpha = 0$, and that each vertex of the triangle is both an origin and a destination. Then, the maximum price of preference is $(1 + \sqrt{3})/2 \approx 1.366$ with one segment, 2 with two segments, and 1 with three segments, all far lower than the theoretical infinite price of preference.

2.4 Obstacles

No-fly zones for drones, such as around airports or over stadiums during events, can be modeled as obstacles in the critical graph. We assume that origins and destinations do not lie within any obstacle, and we replace segments that intersect an obstacle(s) with one or more segments that lie entirely outside all obstacles. For every pair of segments, origins, destinations, and obstacles, the critical graph includes edges between them that meet the following criteria:

- If an edge intersects a segment, it must meet the critical angle conditions in Sec. 2.2.
- If the endpoint of an edge touches an obstacle, the edge must be tangent to the obstacle.
- The interior of an edge must not intersect any obstacle.

Edges between two origins or between two destinations may be omitted, and there may be more segments than in the obstacle-free case if obstacles divide a segment into two or more segments.

How obstacles affect the inefficiency of routing with preferences depends on how they change the optimal and shortest paths, both of which we assume must avoid obstacles. If obstacles only change the shortest path, then preferences have a less detrimental impact on efficiency. If they change the optimal path or both paths, then the outcome depends on how the lengths of the paths change.

Figure 3 shows how obstacles can affect the price of preference relative to the baseline of no obstacle in Fig. 3a. In this example, the upper bold line represents a preferred segment and has a discount factor $\alpha = 0$, i.e., it is free to travel along. When the obstacle increases the length of the shortest path, as in Fig. 3b, then the price of preference decreases. In Fig. 3c, the obstacle increases the length of the optimal path and the price of preference, but the cost of routing around the obstacle in Fig. 3d is high enough that the shortest path becomes the optimal path. Finally, if an obstacle changes the length of both the optimal and the shortest paths, the price of preference may increase, as in Fig. 3e, or decrease, as in Fig. 3f.

(a) Baseline. Optimal path solid via segment (bold), shortest path dashed. Price of preference = 3/2.

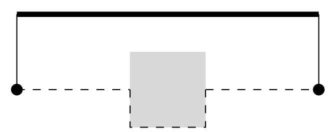

(b) Obstacle that decreases the price of preference to 6/5.

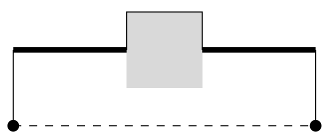

(c) Obstacle that increases the price of preference to 7/4.

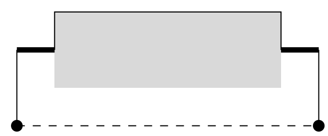

(d) Obstacle that makes the shortest path optimal. Price of preference = 1.

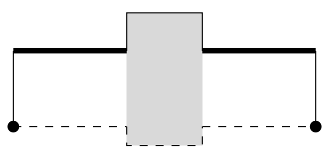

(e) Obstacle that increases the price of preference to 7/4.5 = 14/9.

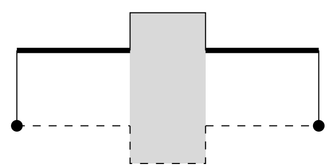

(f) Obstacle that decreases the price of preference to 7/5.

Figure 3: Examples of how obstacles can change the price of preference.

3 Example in a synthetic environment

We first examine how the number of preferred segments and the discount factor affect the price of preference in a synthetic environment.

3.1 Setup

Figure 4 shows an origin, destinations, and segments in a synthetic environment. To understand how the price of preference changes with number of segments, we number them the order in which they are considered in the environment. That is, Segments 1 through n are present when there are n preferred segments.

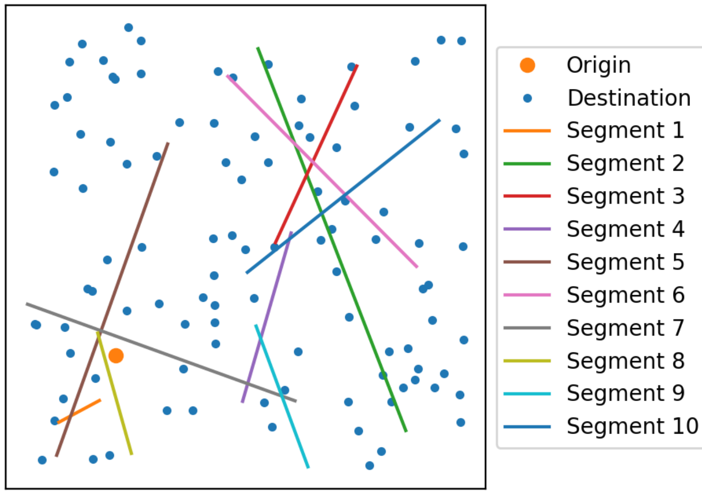

Figure 4: Origin, destination, and segments for a synthetic environment. Colors indicate the order in which segments are added.

3.2 Results

We plot how the maximum price of preference across all destinations changes with the number of segments and the discount factor in Fig. 5. The figure shows how the maximum price of preference across all origin-destination pairs can change drastically with the number of segments. It is not monotonic because with few preferred segments, additional segment are more likely to require a detour to decrease the cost of an optimal path. Once preferred segments reach a sufficient density, however, additional segments may create "shortcuts" that decrease both the cost and the length of the optimal path. In the hypothetical situation where the entire region is preferred, the shortest path is also optimal.

Figure 6 shows the paths corresponding to the sharp changes in maximum price of preference in Fig. 5 for $\alpha = 0$ and 5, 6, and 7 segments. The path with the highest price of preference in each is in bold, and it is not always the same one across different numbers of preferred segments.

4 Drone as first responder (DFR) case study

The City of Detroit Open Data Portal has the locations of Detroit Fire Department fire stations [2] and the incidents that the fire department responded to [3]. In approximately 14% of cases, no incident was found upon arrival at the dispatch address. For these cases, DFR could allow responders to turn around sooner, while for the rest, DFR could improve situational awareness. Thus, we study of how preferences could impact response times.

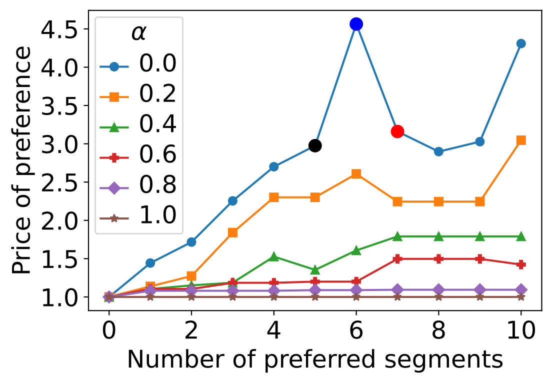

Figure 5: Maximum price of preference for different numbers of segments and discounts.

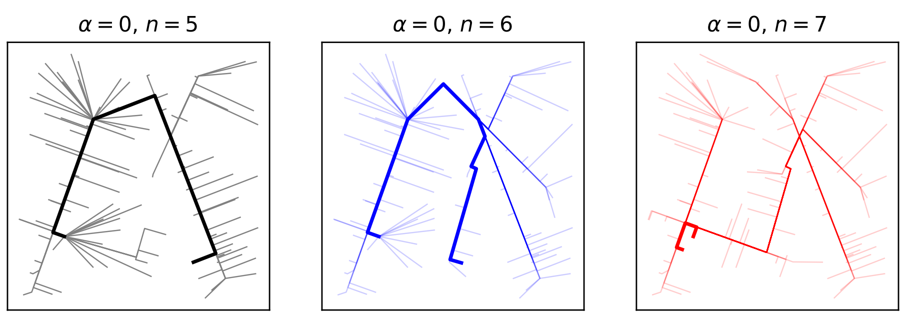

Figure 6: Optimal paths, with highest-price path in bold, for 5, 6, and 7 segments and $\alpha = 0$. Shortest paths are straight lines from origin to destinations.

4.1 Setup

We select five geographically distributed fire stations and 1000 incidents from the year 2024 where there was no incident found upon arrival at the dispatch address. These incidents are intended to be representative dispatch locations rather than specific events, as it is not actually possible to identify which are non-incidents before investigating them.

Because FAA regulations prohibit drone operations over moving vehicles, we model preferred segments along the Detroit River and along several major railways. In contrast to highways and major thoroughfares, these features are typically less heavily used and less densely populated. We also include obstacles to represent no-fly zones over downtown Detroit, where many sports arenas are, and around the Detroit City Airport. This gives us the environment in Fig. 7.

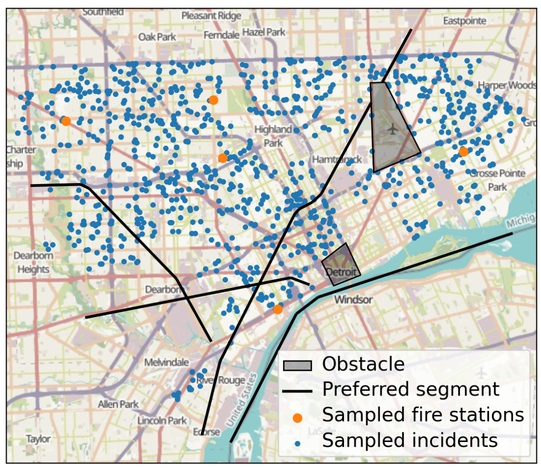

Figure 7: Fire stations (origins) in orange, incident locations (destinations) in blue, obstacles shaded, and preferred segments in black.

4.2 Routing

We find the optimal path from any selected origin to each destination for discount factors of $\alpha = 0$, 0.3, 0.7, and 1. In each of the plots in Fig. 8, the red lines illustrate optimal paths. Blue lines are the shortest paths between incidents and any selected fire station, and these are consistent across discounts. Gray lines link an incident to the fire station that actually responded to it.

There are two complementary things to note as α decreases and the strength of preference increases. Compared to shortest paths, optimal paths concentrate along preferred segments, based on the more saturated red lines along them. As a result, optimal operations cover a smaller area—shortest paths travel where blue is visible, but optimal paths do not. Then, if segments are beneficial to

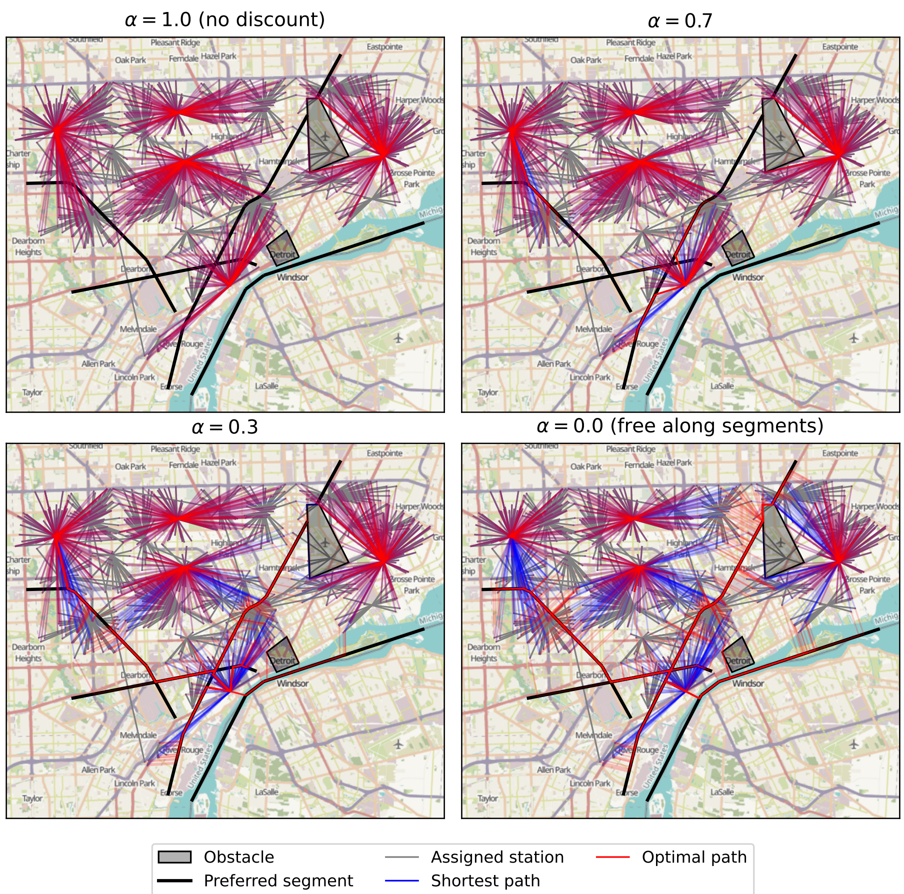

Figure 8: Optimal and shortest paths for different discounts.

safety, privacy, and noise as intended, we can reduce those risks by preferentially routing along them.

Unlike the synthetic example, this case study has multiple origins, and we can observe a radial "star" pattern at each one. These represent incidents for which the optimal path matches the shortest one, a straight line to the nearest selected fire station. Additionally, in this environment, all paths that use *any* segment will have the same origin when $\alpha = 0$, but this is not generally true.

4.3 Results

By comparing the path for each destination when $\alpha < 1$ to its path when $\alpha = 1$, we can find the distribution of the price of preference and of cost reductions due to preferences. For each value of α, Fig. 9 shows the price of preference and the relative cost of each incident. The extents of the shaded regions in the figure indicate the theoretical worst-case price of preference and the theoretical best-case relative cost with preferences for each discount factor.

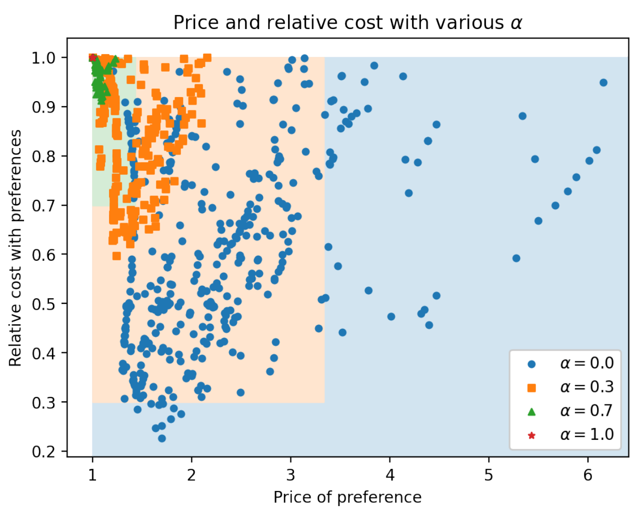

Figure 9: Distribution of prices and relative costs of preferences. Shaded regions indicate theoretical bounds on price and relative cost.

The lower left corner of the figure is ideal. There, we can reap the benefits of preferences—lower safety, privacy, and noise risk—without sacrificing too much efficiency in the form of added travel distance. As might be expected, however, stronger preferences typically lead to lower relative costs but also higher prices of preference. Where a specific origin-destination pair falls in this figure depends heavily on their locations and the environment.

4.4 Travel time comparison

To be used, DFR must be sufficiently timely to complement a ground response. While the routes of ground-based emergency vehicles are not available, response time comparisons are still useful and possibly more relevant.

Figure 10 plots the distribution of response times of a drone with a speed of 20 m/s (72 kph, or approximately 45 mph), which is representative of DFR aircraft, across various α. For comparison, we also include the distribution of actual response times. As α decreases, drones increasingly prefer to travel along segments, resulting in longer response times, but DFR is highly competitive with ground vehicles with the exception of the extremely strong preference of $\alpha = 0$. However, Fig. 10 only shows the *distributions* of response times. We must also ask how using a drone changes response times for each incident.

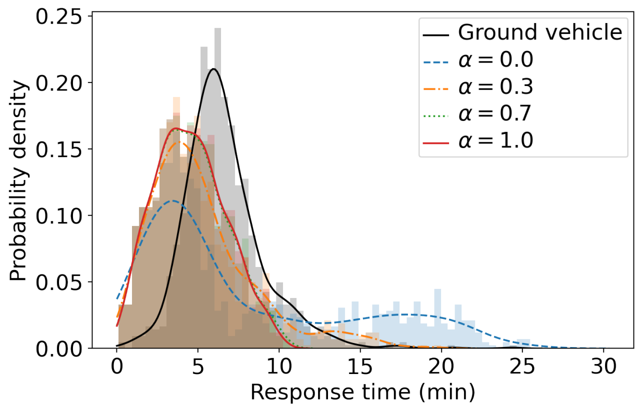

Figure 10: Comparison between theoretical response times of a drone at 20 m/s (72 kph, ∼45 mph) and actual ground vehicle response times.

The distributions of the change in travel time between a drone and that of the corresponding ground vehicle for each α are shown in Fig. 11. Consistent with the distribution of response times, stronger preferences also decrease the response time advantage of drones, but this plot emphasizes that the benefit (or lack thereof) of DFR becomes more differentiated with stronger preferences.

Table 1 lists the median change in response time for the four values of α examined. (Note that this is different from the change in median response time.) While the absolute changes in travel time may seem small, the relative changes are significant in the context of emergency response, where every minute matters and the median actual response time is 6.2 minutes.

As the median change in response time is less than 0 across all values of α, at least half of all incidents would have seen a faster response time with drones than with ground vehicles. On the other hand, there are still many incidents for which a drone would be slower than a ground vehicle, corresponding to positive changes in response time. The share of such incidents is listed in the last column of Table 1.

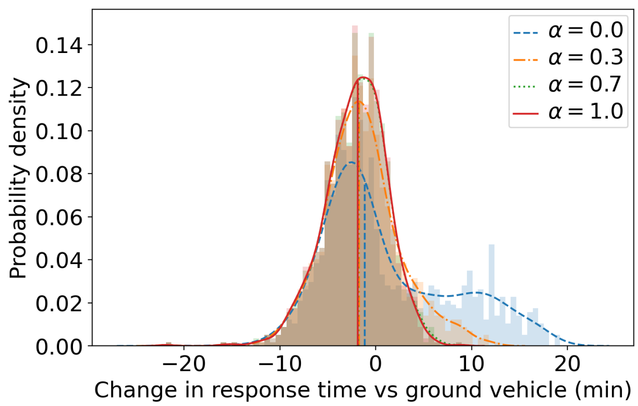

Figure 11: Difference in drone response time compared to ground vehicles for various α. Vertical lines indicate medians.

α	Median change (min)	Median change (%)	Share with slower response than ground vehicle (%)
0.0	-1.1	-18.3	40.9
0.3	-1.7	-27.7	31.0
0.7	-1.8	-29.6	26.7
1.0	-1.9	-30.2	26.6

Table 1: Median change in response time compared to ground vehicle for various values of α.

The results in this section should be considered in the context of having only selected five of the 39 possible fire stations as drone facilities. If more fire stations were equipped to support drones, the advantage of DFR would be even greater. While promising, however, any conclusions should primarily be used to compare across discounts, as this analysis has several shortcomings.

5 Limitations and stakeholder implications

Routing drones with preferences along certain paths is not too time-consuming or inefficient, but our results and analysis should be interpreted alongside their limitations. We also offer some considerations for operators and regulators.

5.1 Limitations

We focus on routing and ignore operational factors, assuming that there is always a sufficiently charged vehicle available for all operations at all facilities. In practice, an operator would have to carefully manage its aircraft fleet. As there is currently still a limited deployment of drones for deliveries or for emergency response, however, the fleet size and supporting infrastructure necessary to meet service or performance targets is unclear.

The choice of how many and which fire stations to equip with drones was arbitrary, and this happened to yield comparable response times to historical trends given our other assumptions. However, these choices could vary, and a facility location problem optimizing those parameters would be relevant in an applied setting. Our conclusions are also sensitive to other environmental factors such as the distribution of destinations, the configuration of preferred segments and obstacles, and the strength of preferences α.

5.2 Implications for operators

The degree to which the price of preference matters will vary by operator. First responders may be reluctant to sacrifice any response time to satisfy preferences, and they must also determine when increased situational awareness from having a drone on-scene before a human team arrives warrants the drone's deployment. This is likely most valuable when the expected travel time to the incident location is significantly longer on the ground than by air due to, for example, distance, congestion, or terrain. Moreover, because drones and first responders on the ground have different capabilities, a team must still be dispatched, at least until a non-incident is confirmed.

On the other hand, with the right incentives, preferences may fly with package delivery operators. Then, the price of preference can be used as a direct estimate of increased delivery times and energy consumption. Even if flights are longer, the time savings of drone delivery compared to ground couriers may still be greater than is suggested in Sec. 4.4 because ground couriers do not have the same priority in traffic that emergency vehicles do.

5.3 Implications for policymakers

Our results also offer insights for regulators considering implementing policies that would incentivize or require operators to follow these preferences.

First is the price of preference itself, as deviating from the shortest path to the optimal one increases distance traveled, corresponding to an increase in travel time and energy usage. Policymakers must consider the trade-off between these factors and their intended goals (e.g., safety, privacy, or quiet) for implementing preferences in the first place.

It is also possible for the reality of routing with preferences to counteract its intended goals. For example, in Fig. 8, the paths without preferred segments are necessarily dense around origins (fire stations) and, to a lesser extent, around obstacles (no-fly zones). Flights also cluster as α decreases, with higher volumes in the smaller visibly red areas at lower discounts. While this is by design for preferred segments, this could also affect edges connecting to or between segments. The result is a potential concentration and exacerbation of the negative externalities that regulation had been intended to address.

Finally, regulators must ensure operators actually follow preferred segments to the maximum extent practical. Whether operators should be rewarded for doing so or penalized for not doing so is an open question. While any incentive structure would likely require external oversight or surveillance to verify compliance with preferences, those may lead to privacy concerns for operators.

6 Next steps

We show how the effects of preferred segments and obstacles on the price of preference depend on their geometric arrangement. Adding a segment or an obstacle can increase or decrease the price of preference depending on how it changes the optimal and shortest paths. In a case study of a drone as first responding application using data from the City of Detroit, we find that (1) strong preferences increase the price of preference, (2) the price of preference rarely approaches the theoretical maximum in practice, and (3) UAVs can have competitive response times for first responder applications even if deployed from a small subset of locations. However preferences are modeled, they must factor into the complex sociotechnical system that arises from drone operations at scale.

Acknowledgements

This work was supported in part through a partnership between Michigan Central and the MIT Mobility Initiative. The authors are grateful to Michigan Central, especially Matthew Whitaker, for insightful feedback and discussions.

References

[1] Burchan Aydin. “Public acceptance of drones: Knowledge, attitudes, and practice”. In: *Technology in society* 59 (2019), p. 101180.

[2] City of Detroit. *DFD Fire Station Locations*. https://data.detroitmi.gov/datasets/b3e02efde01a49fd9f29fb2bdb723000_0/explore. Accessed: June 23, 2025. Feb. 2025.

[3] City of Detroit. *Fire Incidents*. https://data.detroitmi.gov/datasets/d157b0f3decd4c8ea1b0e12c82657552_0/explore. Accessed: June 23, 2025. June 2025.

[4] Roger Clarke. “The regulation of civilian drones’ impacts on behavioural privacy”. In: *Computer Law & Security Review* 30.3 (2014), pp. 286–305.

[5] Reece A Clothier et al. “Risk perception and the public acceptance of drones”. In: *Risk analysis* 35.6 (2015), pp. 1167–1183.

[6] Federal Aviation Administration. *14 CFR Part 107 – Small Unmanned Aircraft Systems*. 2021.

[7] Rachel L Finn and David Wright. “Unmanned aircraft systems: Surveillance, ethics and privacy in civil applications”. In: *Computer Law & Security Review* 28.2 (2012), pp. 184–194.

[8] Zhenyu Gao et al. “Noise-aware and equitable urban air traffic management: An optimization approach”. In: *Transportation Research Part C: Emerging Technologies* 165 (2024), p. 104740.

[9] John Kennedy, Simone Garruccio, and Kai Cussen. “Modelling and mitigation of drone noise”. In: *Vibroengineering Procedia* 37 (2021), pp. 60–65.

[10] Dasom Lee, David J Hess, and Michiel A Heldeweg. “Safety and privacy regulations for unmanned aerial vehicles: A multiple comparative analysis”. In: *Technology in Society* 71 (2022), p. 102079.

[11] Joseph SB Mitchell. *Shortest paths among obstacles, zero-cost regions, and roads*. Tech. rep. Cornell University Operations Research and Industrial Engineering, 1987.

[12] Beat Schäffer et al. “Drone noise emission characteristics and noise effects on humans—a systematic review”. In: *International journal of environmental research and public health* 18.11 (2021), p. 5940.

[13] Qichen Tan et al. “Exploring noise reduction strategies: Optimizing drone station placement for last-mile delivery”. In: *Transportation Research Part D: Transport and Environment* 133 (2024), p. 104306.

[14] John Villasenor. “Observations from above: unmanned aircraft systems and privacy”. In: *Harv. JL & Pub. Pol’y* 36 (2013), p. 457.

[15] Roland Weibel and R. John Hansman. "Safety Considerations for Operation of Different Classes of UAVs in the NAS". In: *AIAA 4th Aviation Technology, Integration and Operations (ATIO) Forum.* Reston, VA, USA: American Institute of Aeronautics and Astronautics, 2004.