

# Alpha Pilot Challenge Webinar Q & A Log

## General Questions:

1. What is the goal of AlphaPilot? Is it recruiting?
  - a. AlphaPilot is about building community and trust in autonomy. Our goal is to bring together users from all walks of life and create a fun, common challenge that pushes the state-of-the-art.
2. Is there any provision for under 18-year-old team members?
  - a. Unfortunately, team members must be the age of majority in their country or residence or at least 18 years of age, whichever is older as of the start of the Challenge. You can view the complete details on eligibility in the [legal agreement](#).
3. A portion of each test is based on the judges' evaluation of our technical report. What are you looking for in these reports?
  - a. There are a few keys things we are looking for in the technical report: teams to explain clearly what they did, how their approach would (or would not) be used in AlphaPilot, and any challenges they ran into how and how they overcame them.
4. How will you use our source code and intellectual property?
  - a. Lockheed Martin, HeroX, the Drone Racing League, and NVIDIA (collectively, "AlphaPilot Administration") does not claim intellectual property rights to any code submitted during the course of AlphaPilot. Please see the legal agreement and the FAQ for more information
  - b. Qualifier Source Code: Source code was requested during Qualifiers to assist AlphaPilot Judges in selecting Finalist teams. All code will be hosted and secured solely on HeroX servers, accessible only through the HeroX platform. AlphaPilot Judges will be the sole viewers of source code and will be under a non-disclosure arrangement
5. What expenses will teams be responsible for covering?
  - a. The competition is designed to minimize team expenses. There are no fees or travel requirements for the Qualifier. The only cost that Finalists will incur is travel to AIRR events. Only one member of each Finalist team is required to

attend each AIRR event. Teams are permitted to train however they see fit, including the purchase of additional resources beyond the development kit. However, this is not required and is at the discretion (and expense) of teams.

6. You mentioned teams are encouraged to find funding support. Do you have any existing resources or support that you can share to get connected to funding resources?
  - a. Lockheed Martin and DRL will work with teams to secure sponsorships. This could be from universities, companies, or other organizations. We will leverage our connections and distribute team media to the extent possible.
7. Can teams collaborate at sub system level? We are getting good at software, we could partner with a team already good at flying?
  - a. Yes! Teambuilding is encouraged and can include team mergers. The Team Matching function is there to help. Please solidify any such changed by the end of registration.
8. What GPU architecture should we optimize our code for?
  - a. The following instance will be used at the basis for the testing environment: <https://aws.amazon.com/marketplace/pp/B077GCZ4GR>. Please references the spec here for details such as available RAM.
9. Can we run a docker container on the server?
  - a. Unfortunately, no. The testing procedure is automated and set, and a docker container is not compatible with that.
10. Who are the judges for AlphaPilot?
  - a. The Qualification judges are a mix of experts from academia, industry, and government. Many experts in drones, GNC, machine learning, machine vision, and more support AlphaPilot.

## Test 1

11. What information are you looking for in the Test 1 video? We have a team of 10 people, what should we prioritize given we only have three minutes.
  - a. AlphaPilot is a year-long endeavor. Aside from technical performance in Test #2 and Test #3, we want to know that your team is a solid investment. Do you have the cohesion and commitment to see the challenge through? We don't want teams that will break-up, become distracted, or lose interest half-way through.

Do you appreciate the difficulty of AlphaPilot and the level of effort that will be required? Have you thought through possible setbacks? Can you reference some examples of difficult software problems and how you handled them? What makes your team unique – i.e. more than a group of individuals? In summary, make us feel confident in you as a team.

## Test 2

12. Is there a test instance on your platform we can use to see that our code will work when you score it?

- The starter scripts and scorer scripts provided for Test 2 represent how each team's source code will be tested.
- We will not be providing the exact test instance for teams. However if a team defines a *GenerateFinalDetections()* class with a *predict(self,img)* function that runs smoothly in *generate\_submission.py* (as shown in the starter scripts), this provides the sanity checks needed on source code.

13. How do you recommend we use the ground truth labels in the training dataset?

- Above all, follow the guidance given in the Test 2 description for accurately identifying correct gate labels in images. Some of the training data ground-truth labels do not follow these rules accurately, and that should be considered when developing machine vision algorithms. Handling real-world challenges effectively will be critical for success in the AlphaPilot Competition, and similarly, teams need to deal with some flawed ground-truth in Test 2 and sensor noise, control drift, and modeling errors in Test 3. Address how your team plans to deal with these in the Technical Reports, because AlphaPilot judges would like to know this when considering a team's approach.

14. How do the ground truth labels work in the testing dataset? Is there a "fudge factor" that will score correct if we're a few pixels off?

- The ground-truth (GT) labels represent the 4 coordinates which define a polygon that maximizes the flyable region of the AIRR gates. The primary measure used in the *MAP* score is the Intersection over Union (IoU). The IoU is computed by dividing the area of overlap between the GT and predicted bounding boxes by the area of union. For more information on the implementation of this metric used for AlphaPilot, read more here: <https://arxiv.org/abs/1405.0312> instance here <https://aws.amazon.com/marketplace/pp/B077GCZ4GR>.

15. Is calling a C/C++ library from Python allowed for Test 2? I'm working under assumption that it is allowed, since almost all Python libraries, where performance is an issue are written in C/C++ or other higher performance language.

- This is allowed. However, please make sure your code is still compatible with the specified testing environment and requirements.

16. Can a 3D model of the gate used in Test 2 be made available?

- Unfortunately, we won't be able to provide any mechanical drawings of the gates.

17. Is there any difference in the distribution of the public and private leaderboard data for Test 2?

- No

18. Do you account for inference time?

- The measure of execution time is assessed according to the wall clock time for the *generate\_submission.py* function to run *finalDetector.predict(img)* as defined by your class. That is, the time for the gate detector to read an image, extract the corners of the flyable region of the gate, and output the label.
- The total Algorithm Score is then calculated by subtracting the average wall clock time from the weighted MAP score and multiplying by 35 to get a maximum of 70 points:  $Score_{Algo} = 35 \cdot (2 \cdot MAP - avg\_time)$

19. Do we need to detect gates obscured by pillars?

- Each gate should be labeled with a single polygon with 4 corners around the flyable region (even if it's partially obstructed). When all 4 corners of the flyable region of the gate are visible, teams are expected to find them (even if part of the flyable region is not visible). If that is not the case (where one or more corners is not visible in the image), teams will not be tested on these images, and they have been removed from the testing data. This is largely due to the fact that it is very difficult to accurately label the ground-truth for those cases.

20. Could you elaborate on the MAP score for test 2?

- Please read the paper included in the Test 2 overview. We also suggest teams do a bit of their own research into the score. This is a very common metric for data science competitions, and we utilize a common implementation in AlphaPilot.

21. Execution time for test #2 will be benchmarked on multicore, hyperthreaded CPU where my task execution, due to primarily cache thrashing, may be randomly affected by another task running on that CPU at the same time. Are you aware of that?
- Each team's algorithm will be tested on the server sequentially (one at a time), and no other tasks will be running on the server during that process.
22. Can you expand on the confidence score?
- a. There exist two distinct measurements needed in object detection:
    1. Whether the image is correctly classified (i.e. if an object exists in the image)
    2. How well the object has been localized
  - b. Simple metrics introduce biases and so it is important to assess the risk of misclassifications. Thus, the "confidence score" is used. For more information on how this is used in AlphaPilot, please do some reading into the "mean Average Precision" on the *MAP* metric and how the confidence score is considered mathematically.

## Test 3

23. What are some good resources to learn ROS and learn how to add our nodes onto the flightgoggles?
- a. Please refer to the official ROS tutorials: <http://wiki.ros.org/ROS/Tutorials>
24. I couldn't find the API documentation for FlightGoggles to send http messages over to the simulation. Can I use any ROS API to do that?
- FlightGoggles uses standard ROS message types for communication. As such please refer to the ROS message publisher/subscriber tutorial to familiarize yourself with this message passing concept:  
<http://wiki.ros.org/ROS/Tutorials/WritingPublisherSubscriber%28c%2B%2B%29>
25. Will FlightGoggles be updated to depend on the current recommended ROS version ( --roscat melodic v.s. --roscat kinetic)?
- We decided to start with Kinetic for stability. MIT will update and maintain flightGoggles in the future. We will consider an update to melodic in the future. However, it will happen after test 3 is completed.
26. Are we allowed to change the camera angle in the FlightGoggles simulator? Can you set the camera angle to something greater than 0?
- No, we are not allowing teams to change the camera angle. All teams are challenged to use the same hardware setup in AlphaPilot.

27. What data inputs can be used in Test 3?

- Please see the README file for Test 3 for a complete list of allowed and prohibited inputs.

28. Are the gates disturbances in the IR beacons sensor or in the real position of the gates in map?

- The gate disturbances are from the nominal position of the gates in the map.

29. Will there be a left/right camera version of the IR beacons in Flight Goggles?

- Please see the description provided by alphapilot for allowed topics and params: <https://www.herox.com/alphapilot/resource/320>

30. In Test 3, it's hard to know when we fly through a gate, without ground-truth data. Can we get an event for us to subscribe to?

- Unfortunately, we will not provide this for teams as this is not information you would realistically have during a drone race.

31. Will all target gates be vertical, or will there be horizontal gates (flying through vertically)?

- All target gates in the virtual qualifiers will be vertical for both Test #2 and Test #3.

32. For test 3, will the gates locations (not including the perturbations) be changed for the final test?

- No, the nominal positions of the gates are the same for both the Leaderboard practice and final tests.

33. How will the race be timed? Will the timing start from begin of code launch or will it start when flying through the first gate or leaving a bounding box? (technically: how will time to initialize estimators etc. be penalized?)

- a. The race will be timed from the moment of the drone's first takeoff movement to the time when the drone arrives at the finish gate.

34. From the latest update it would seem that the ROS server is sending a message about the location of each gate. Are the messages sent from the given ROS node providing a x,y from the camera view from the drone's perspective?

- The `"/uav/camera/left/ir_beacons"` ROS topic publishes all *unoccluded* IR beacons within the camera's field of view. This message contains IR marker detections in the image space of the left camera.

35. What are minimum requirements for FlightGoggles video card wise?

- Minimum Local Hardware Requirements
  - A Vulkan-supported GPU with  $\geq 2.1$ GB of VRAM.
  - Native Ubuntu 16.04 installation with ROS Kinetic on x86/x64.
  - Virtual machines do not have native access to the GPU and therefore are not supported.
  - nvidia-docker2 does not currently support the Vulkan API and therefore cannot run the FlightGoggles renderer binary (see workarounds in issue #46).
- We have tested this project on two different setups: High end Desktop computer with:
  - Processor: Intel i9 extreme (i9-7980XE)
  - RAM: 32Gb
  - GPU: Titan V
- We have also tested on the following AWS instances:
  - p3.2xlarge
  - g3s.xlarge
- a. Please see <https://github.com/mit-fast/FlightGoggles/wiki/Prerequisites-and-Testing-Setup>

36. Does the flightgoggles simulator take .yaml files as an input or will we have to convert the files?

- Teams can use the Challenge YAML files as input to FlightGoggles for their own development and testing.
- a. However, users of the FlightGoggles simulator should not have to read or modify any YAML files for correct operation of the simulation system.
- b. All relevant information to the FlightGoggles simulator is exposed through the ROS Parameter API: <https://github.com/mit-fast/FlightGoggles/wiki/ros-params>

37. Can you simulate more than one drone?

- We intentionally did not put support for multi-drone simulation in order to avoid confusion and to keep the initial FlightGoggles code as simple as possible. Before the challenge is done, we will not provide the software for simulating multiple

drones. That said, you are able to run multiple instances of FlightGoggles if you would like to conduct several tests, of course.

38. Keras is now a part of tensorflow. Can we use keras to develop the AI model.
- Solutions need to be compatible with the algorithm and testing requirements. Otherwise, there are no restrictions on the approach and libraries teams can use to develop their solutions.

## Post-Qualifier Round

39. Will camera distortion or field of view change in the competition?
- a. Yes. The AlphaPilot Competition will use different hardware than in the qualifications.
40. Will sensor noise levels be the same? (so that estimators etc. can be roughly tuned to it)
- b. No. The AlphaPilot Competition will use different hardware than in the qualifications.
41. Will the teams who do not win the Qualifier Round be able to access DRL's simulator in the future?
- c. Although DRL doesn't currently have plans to support a public flight autonomy simulator, this option is still being considered
42. Teams will be provided a drone for the races in the fall. Can teams develop their own autonomous drone?
- d. To ensure vital features and functionality are included and a level competition field is provided for all participants, all teams will utilize the autonomous drone build provided by DRL.
  - e. That being said, teams are encouraged to use whatever additional resources they have available to them for development and testing ahead of the races.
43. Unlike software, hardware often perishes in bugs and crashes, and parts are needed for replenishment, so are there any spares, replacements available? How does this work for international participants? Shipping, customs etc.?
- f. All drones will be provided, maintained, and repaired by DRL.
44. Since the Jetson Xavier modules with devkit board alone weights close to 700g, can you share the specs for the drone platform that will be used to carry it?
- g. Yes, this will be shared with the final qualifying teams



45. How many races will there be and the potential locations?

- h. There will be four locations in 2019. The final locations have not been selected but are all planned to be in the US.

46. Will the kit include a drone/parts to test with? Battery management and drone hardware is part of racing and working with live hardware I would imagine is important to training for the race

- i. Yes, final selected teams will receive supplementary hardware assets for testing along with their development kits.

47. If I don't qualify in the virtual challenge, can I still participate in the races?

- j. We are currently exploring whether this opportunity will be made available to teams.

48. Will we be able to modify/customize our drones for just aesthetically purposes?

- k. It's important that all drones are exactly identical in form and function so we keep a strict policy on no modifications.

49. Will teams be responsible for maintaining their drones or will there be a team to help to keep each consistent?

- l. All drones will be provided, maintained, and repaired by DRL

50. Will gates' markings or printing on them be same in actual race as in test 2 dataset?

- m. Yes.