

### Short Description:

A team of qualified, experienced industry professionals and researchers with backgrounds in AI/ML/RL, sensors, robotics, hardware and software eager to compete and win.

Please provide the name and affiliation of each team member. Please designate the team captain with an asterisk. \*

\*Brian Lee - General Motors, Georgia Institute of Technology  
Vinayak (Vin) Kankanwadi – Tabcorp, Brisbane  
Ian Liu - West Monroe Partners, Georgia Institute of Technology  
Steven Royster - West Monroe Partners  
Marcus Schwarting - Argonne National Laboratory, Georgia Institute of Technology  
Ziad Abou Wasa – Baker Hughes, Georgia Institute of Technology  
Sanjeevkumar Umadi – Tata Elxsi  
Sagar Eligar - Autonomous Underwater Vehicle Team, RIG-NITC  
Brendan Walsh - Former Product Manager, Cohen & Steers

Should your team qualify, please share a 500 character description of your team to use in the announcement? \*

We are a team of seasoned engineers in autonomous domain, software engineers, and research scientists who are dedicated to complete our design, follow through with integration, and aim to win the AlphaPilot competition. SkyNet is our name, autonomous drone racing is our game. Just like the skynet in the movies, our model will be intelligent in that it will learn the best way to fly over time with the help of reinforcement learning techniques, but will also mix the right amount of conventional approaches where RL is intractable to apply.

Is your team on social media? Do you have a website? Please share those links here.

The LinkedIn pages for each team member are as follows:

Vinayak (Vin): <https://www.linkedin.com/in/vinayak-kankanwadi-a956ba44>

Marcus: <https://www.linkedin.com/in/marcus-s-a7779486/>

Brian: <https://www.linkedin.com/in/brianhanullee/>

Ian: <https://www.linkedin.com/in/ian-liu/>

Steven: <https://www.linkedin.com/in/steven-a-royster/>

Ziad: <https://www.linkedin.com/in/ziad-abou-wasaa-a075743b/>

Sanjeevkumar: <https://www.linkedin.com/in/sanjeevkumarumadi/>

Sagar: <https://www.linkedin.com/in/sagar-eligar-9306b816b/>

Brendan: <https://www.linkedin.com/in/dbrendanwalsh/>

Team video: <https://youtu.be/Xt9x3oGh8gw>

Where is your team located? \*

Various locations including New York City, NY; Chicago, IL; Houston, TX; and Sterling Heights, MI; Brisbane, Australia; Kerala, India; Karnataka, India.

1. What are the big technical challenges an autonomous drone needs to overcome in order to beat a human pilot flying the same drone? Why? \*

To better discuss some of the technical challenges, we will use a somewhat analogous example of driverless vehicles. Driverless vehicle technology is an area of expertise for some of our teammates. The technical challenges, while being different with regards to controls in 3D space, are similar in nature.

Driverless vehicles outperform human drivers in a number of regards. The most obvious of these is that a human has a limited attention span and perception system. Humans suffer from slow reaction times, fear, and will often misjudge complex situations. Computers, however, do not suffer from such deficiencies. A quality sensor suite, including lidar and radar detection, monocular and stereo cameras, and other sensors can likely outperform the abilities of human attention and perception. Some of these techniques also have advantages in areas where humans are naturally lacking; for example, radars and lidars can work well even in low-light scenarios.

The challenge of utilizing computers and sensors for driverless vehicles therefore becomes a problem of limited computational ability for on-the-fly calculations. Computational efficiency becomes absolutely critical. Optimization of the entire software stack and its design, from sensing and localization to path and behavior planning for control, becomes an important goal for driverless vehicles.

Many of the same deficiencies that humans suffer for driverless vehicles are known issues for drone racing. On top of that, the biggest challenge for an autonomous racing drone is operating under the "acro" flight mode. This mode of flying is used by professional drone racers as it allows for the most fine-tuned responsiveness and the highest level of control. However, it implies all "self-leveling" corrective features be turned off. This means every single input from the pilot is reflected in the drones flight control system. It allows for flips, rolls, and advanced custom PID tuning. FPV professionals and enthusiasts often describe "Acro" mode as a flight mode that requires constant "course-correction." Once a stick input is received, the drone follows that flight course, and subsequent stick movements become a kind of constant "course correction."

Autonomous systems for drone racing will face similar obstacles as autonomous vehicles regarding stack optimization, with the additional challenges of "Acro" mode. One of the greatest challenges for an autonomous drone system will be establishing an effective relationship between localization and control in "Acro" mode. A human can rapidly localize a drone with high accuracy for decision-making, while this is far more computationally complex for a machine. When lidar and radar are available, localization accuracy would quickly become more tractable, but given only camera inputs, localization accuracy becomes solely dependent on detection data for visual odometry and inertial motion units (IMUs). Racing drones are optimized for weight, and fly at speeds in excess of 100mph with hairpin turns and significant changes in acceleration and deceleration.

For both driverless vehicles and autonomous drone racing, control will take place using similar techniques of principle-integrated-derivative control (PID) and model predictive control (MPC). The task becomes more complex for autonomous drone racing since a drone has four necessary

axes of control (thrust, roll, pitch, yaw) while a driverless vehicle has only two necessary axes of control (thrust and yaw). Ultimately methods incorporating reinforcement learning will be utilized to enhance performance, particularly in cases where unexpected environment noises (such as sudden unknown winds) will be experienced by the drone.

2. Describe your team's planned technical approach to AlphaPilot. Why do you think your team could win with this approach? \*

For autonomous systems, the main challenges include perception, mapping, localization, behaviour and path planning, and controls subsystems. There are a number of methodologies available to tackle these problems. For the perception required in Test 2, we chose to implement two algorithms, then determine which would perform better based on the given scoring function. We chose to implement "You Only Look Once" (YOLO) and Mask R-CNN. In the end, we determined that Mask R-CNN was more accurate with sufficient inference time.

While our solution is already running quickly on an AWS instance, additional improvements will be made to attain higher frame rates suited to the faster speeds of racing drones. More tailored training in addition to data augmentation could increase the accuracy significantly, while further optimization using CUDA parallelization or Tensor-RT optimization will further speed up the computation time of our implementation. We believe that a very accurate detection in terms of IoU isn't necessary to win the race. Although the IoU should be sufficiently accurate, it certainly isn't required to find enough of the flyable region of subsequent gates to plan the path and fly through them. We understand the tradeoff between speed and accuracy in these types of robotic systems, and we believe a high-frame-rate solution with good enough IoU detection is the best perception system for this type of drone racing.

For localization, we are planning to utilize various kinds of Bayesian approaches, such as graph-SLAM and/or extended Kalman Filters (EKFs). Inertial motion units (IMU) and visual odometry are part of our feedbacks as the states, the location of our drone will be constantly estimated. While utilizing the known locations of the gates, whose 3D locations are estimated from perception system, we aim to improve our localization accuracy and prevent bias offsets that are known to occur for IMUs.

For path planning, since the race course and order of gates are known, a simple path 3D spline planning can be utilized for test-3. For the real race, path planning algorithms such as A\* and Dijkstra will be utilized to plan a path that optimally avoids obstacles in real time. In complicated situations where other drones are nearby, reinforcement learning will be explored to assess whether it can plan paths that maximize speed while avoiding collisions better than conventional A\* or Dijkstra approaches.

The greatest challenge will come with controlling the drone. This is where groups will likely diverge in approaches. Our team will start by implementing a simple principle-integral-derivative (PID) control for the system, then move to a more complex approach of model-predictive control (MPC). However, eventually we intend develop a reinforcement learning approach in order to account for all unknown environment noises that

can occur. This could be particularly useful when our team is ready to move away from the simulation and into real-world testing. Various perturbations due to wind, drift, and other forces could be more precisely accounted for by a reinforcement learning method than a traditional PID/MPC architecture. No PID or MPC can account for future unknown noises better than reinforcement learning. However, PID and MPC will be developed so that we can compare with the RL's performance and make the final decision to see which one will fair better for the race.

Our team believes that in developing this reinforcement learning approach, our flights may be more robust to challenges in the course as well as less susceptible to outside perturbations that a PID may not handle gracefully.

3. How does your team plan to use the simulator and development kit provided by DRL? How does your team plan to handle real-world variations that are difficult to capture in simulation? \*

Simulators are very important in the autonomous domain since real-world testing is often scarce due to the difficulty of data collection and aggregation. Simulations will allow us to try a number of scenarios that may be difficult to replicate with real-world conditions. Simulations can also allow users to generate custom conditions to challenge the robustness of a method (for example, seeing how a drone might perform if initiated without a direct sight of a nearby gate). Simulations can be also easily parallelized (multiple simulations at once) if necessary. Training using reinforcement learning techniques for behaviour planning can require an enormous number of iterations over many episodes. Training on this scale would quickly become intractable without a simulation.

While simulators are fantastic for training models, there is no substitute for testing in real-life scenarios. In some cases, a simulator may be lacking in its interpretation of real-world physics and kinematics, which would catastrophically impact the drone performance. There are also a lot of real-world effects such as lighting consistency, perception distortions due to rapid movements, and large perturbations which can be attributed to even mildly windy conditions.

There are several ways that one could overcome the challenges introduced by moving from “in silico” to “in situ”. The first and simplest change to implement would be a Bayesian approach to account for covariances from multiple sensors due to noise and other effects. To deal with wind perturbations, one could manually fine-tune PIDs or other control techniques multiple times such that even against wind perturbations, the drone will maintain its course. But we believe that the correct solution for real-world situations is reinforcement learning algorithms, which can automatically account for these errors if properly trained.

4. What support, resources, and tools does your team plan to use for the competition to supplement the hardware, software, and training provided? How do you plan to support necessary travel should you proceed in the competition? Please include any hardware, software, people, data, mentoring, sponsorships, etc. \*

Hands-On Drone Experience and Other Hardware:

One of our team members (Vin) has experience constructing quad- and hexacopter drones from carbon fiber, and subsequently writing custom software to fly them. Additionally, several of our members own drones. Two of our members (Ian and Brendan) own and regularly fly custom-built drones, including a 7-inch Jetson TX2 with a PX4 flight control system that we intend to use for real-life testing after we qualify. We are looking forward to running our custom software on these drones with the target platform mounted and comparing them against our experienced human pilots.

Many of our members have powerful desktop computers (GTX1080s and higher) and high performance computing servers (100xTesla P100) at their disposal for heavy DNN or RL training.

#### School Experience:

Four of our members (Brian, Ziad, Marcus, Ian) are enrolled in the online master's degree in computer science (OMSCS) at Georgia Institute of Technology. All four members have denoted a specialization in machine learning and artificial intelligence. Classes offered include those taught by Dr. Michael Littman, Dr. Charles Isbell, Dr. Sebastian Thrun, and other noted academicians and innovators. And our Teaching assistants who work in the industry applying reinforcement learning to flying agents, have agreed to help with our RL related queries should they arise.

#### Work Experience:

One of our members (Vin) has worked on contracts with Lockheed Martin in the past, and currently operates within the defense sector. Another (Brian) works for GM in their autonomous driving initiative. And another (Marcus) is an AI research scientist who build models for materials science applications at Argonne National Laboratory.

#### Regarding Travel:

All members located within the continental United States (Brian, Marcus, Ian, Steven, Ziad) are willing and able to travel to events related to AlphaPilot. One member (Brian) would ask for travel sponsorship from General Motors while another member (Marcus) would ask for travel sponsorship from Argonne National Laboratory. The other two members have expressed willingness to pay out-of-pocket to attend events. All other members will be available to connect via phone or video chat.

5. What do you see as the biggest challenges specific to your team that you will face during the AlphaPilot competition? What is your initial plan for overcoming this challenge? For example: missing skills, financial challenges, team cohesiveness, time commitment, etc. \*

For many members, the time commitment may be an issue from time to time. Most of our members have full-time jobs in highly demanding industries. Around half of us are still working on furthering our education. While this can be problematic, the problems that most of us are engaging with at work and in school are building up our knowledge base and making us more adept at contributing positively to AlphaPilot. For example, three of our members are currently enrolled in a reinforcement learning class taught by Dr. Michael Littman. The knowledge we gain from this class will be readily applied to drone control for Test 3 and beyond. Other members have taken classes in machine learning, artificial intelligence for robotics, and computer vision.

In addition to other time commitments, some members lack experience and knowledge on various technologies required by the competition. Two of our members (Brian and Vin) have practical experience implementing all individual parts required to complete both Test 2 and Test 3, and can help guide other team members in the right direction by providing learning materials and past projects.

Other members are able to contribute in smaller but important roles while working in their area of specialty. For example, one team member (Ziad) has a strong background in control theory due to a background in electrical engineering. With very little assistance, he has made significant strides towards development of a working PID for Test 3.

Summarize the key points you shared in your video in response to the question: Describe your team's composition, background, and qualifications as they relate to the AlphaPilot challenge and associated technologies. Why is your team the one to beat? \*

We have members who are highly experienced in autonomous systems such as self-driving car and autonomous drones, as well as passionate professionals who are taking advanced computer science classes together to gain further understanding, all while juggling full-time work. Two members (Brian and Vin) have created successful designs for the integration of complicated subsystems involving both hardware and software. These past designs are considered similar in scope and magnitude to the problem given in this challenge, and can yield correct guidance to resolving the complex problems that will arise along the way. While taking on these problems can be challenging and time-consuming, past experience with similar challenges shows our fortitude and persistence in striving towards our goals, no matter what. Add to this a diverse set of individuals specializing in related areas of development, we believe we possess the right set of people and skills to compete and win this challenge.

Summarize the key points you shared in your video in response to the question: Have any team members participated in prior competitions focused on artificial intelligence and autonomy? Has your team worked together in the past on technical challenges with comparable scope or complexity? Please give examples and describe the results. \*

Vin - AWS DeepLearning Brisbane (2nd place): Model to identify Ships  
Sagar - participated in Singapore Autonomous Underwater Vehicle 2018 (SAUVC2018). The purpose was an autonomous bot to be used for cleaning swimming pools.  
<https://www.youtube.com/watch?v=7rnf42z0WM&feature=youtu.be>

Brian - Burningman competition: built 35-foot long robotic snake called titanoboa (world record for the largest robotic snake), which made a splash in the media.  
<https://www.youtube.com/watch?v=UZk6bdwudDg>

Also, while not really a competition like AlphaPilot, most OEM car manufacturers (such as GM where Brian is part of) and software heavy companies like Waymo and Uber are all in this giant competition of who can get "there" as fast as possible with the most robust and efficient

solution. While interacting daily with industry leaders and developers working on autonomous driving technology, Brian is used to fast-changing and fast-moving pace where various technologies get developed, integrated, and tested. Supercruise product was well received by the media and the public.

Brian, Ian, and Marcus have worked together in solving reinforcement learning problems in the class they are taking together at Georgia Tech as part of Computer Science Masters Program.