


The University of Oklahoma

Sooner Competitive Robotics
"Suspended Disbelief"

Submitted May 15th, 2026



Competitions:

AutoNav

Self Drive

Name	Email	Role
Kyle Engel	kyle.engel@ou.edu	Team Captain
Dylan Zemlin	dylan.zemlin@ou.edu	Team Mentor
Noah Zemlin	noah.zemlin@ou.edu	Advisor

I certify that the design, development, and work put towards this project by the Sooner Competitive Robotics students is significant and equivalent to a senior design course.

X *Noah Zemlin*

33rd Annual Intelligent Ground Vehicle Competition (IGVC)
May 29th – June 1st, 2026
Oakland University, Rochester, Michigan

Table of Contents

1	System and Subsystem Requirements	1
1.1	Systems Engineering Process	1
1.2	Requirements	1
2	Mechanical Design	3
2.1	Payload & Weatherproofing	4
2.2	Requirement Traceability and Measurements	5
3	Safety	6
3.1	Requirement Traceability and Measurements	7
4	Electrical and Electronic Design	8
4.1	System Overview	8
4.2	Computing Platform	8
4.3	Custom PCBs	9
4.4	Requirement Traceability and Measurements	9
5	Perception	10
5.1	AutoNav Perception	10
5.2	Self Drive Perception	11
5.3	Requirement Traceability and Measurements	12
6	Driving Logic	13
6.1	AutoNav Driving Logic	13
6.2	Self Drive Driving Logic	14
6.3	Requirement Traceability and Measurements	15
7	Key Performance Indicators	16
7.1	Overview	16
7.2	AutoNav Key Performance Indicators	16
7.3	Self Drive Key Performance Indicators	16
8	Analysis of the Complete Vehicle	17
8.1	Lessons Learned During Construction and Integration	17
8.2	Software Testing, Bug Tracking, and Version Control	17
8.3	Simulation Based Testing	17
8.4	Physical Testing	17
8.5	Overview	18
8.6	Vulnerability Summary	18

Section 1: System and Subsystem Requirements

1.1 Systems Engineering Process

Our team derived system and subsystem requirements through a top down and use case analysis of the IGVC 2026 rules.

For the top down analysis, we started at the mission level: the robot must navigate autonomously, avoid obstacles, follow lanes and waypoints, and respond to operator commands. We broke that down into five sections: drivetrain, perception, autonomy/driving logic, electrical/power, and safety.

For the use case analysis, we walked through scenarios the robot would face such as lane following, obstacle avoidance, waypoint navigation, estop activation, battery depletion, and others. Each exposed additional requirements around localization accuracy, estop latency, and stability over rough terrain.

1.2 Requirements

Mechanical

Name	Driver	Measure	Target
Operating Speed	Rules cap all runs at 5 mph; drivetrain must enforce this mechanically and electrically.	Straight-line run at max speed; compute speed over measured distance.	≤ 5 mph
Ramp Traversal	Course includes ramps; a stall ends the run.	Drive up increasing inclines; record max slope handled.	Capable
Structural Stability	Camera-first stack means frame flex directly hurts perception.	Drive rough terrain; observe camera shake.	Visibly stable

Safety

Name	Driver	Measure	Target
Wireless E-Stop Latency	Required by rules; delay between press and stop is a direct safety risk.	Max-speed run; measure time and distance to full stop after estop press.	≤ 1 s
Physical E-Stop Access	Rules require a visible, reachable physical estop.	Verify line of sight, accessibility, and function.	Rule height; stops robot.
Safety Light Visibility	Rules require a flashing indicator during autonomous operation.	Test outdoors in full sun.	Visible flashing in sun.

Electrical / Electronic

Name	Driver	Measure	Target
Battery Voltage Under Load	8 motors, Jetson, cameras, and PCBs off one pack; can't brownout mid-run.	Log terminal voltage over a full peak-load run.	$\geq 11.0\text{ V}$
EStop Signal Integrity	EStop line shared across all PCBs; must be driven only by the designated board.	Measure signal in both estop and normal states.	3.3 V normal; low on estop.

AutoNav Challenge

Name	Driver	Measure	Target
Lane Line Detection Rate	Lane departures are penalized; cameras must reliably detect white lines outdoors.	20 m segment replay; fraction of frames with both lines correctly detected.	$\geq 95\%$
Obstacle Detection Range	At 5 mph we need 3 m of detection range to plan and execute avoidance in time.	Cone at measured distances; min range for $\geq 80\%$ confidence across 20 trials.	$\geq 3\text{ m}$, $\geq 90\%$
Lane-Center Tracking Error	Centerline margin determines boundary violation risk on curves.	Lateral offset at 10 Hz from camera midpoint over 20 m.	$\text{MAE} \leq 15\text{ cm}$
Course Completion Rate	Scoring is based on penalty-free completions.	Count clean laps across all competition runs.	≥ 2

Self Drive Challenge

Name	Driver	Measure	Target
GPS/IMU Localization Accuracy	Self Drive navigation is lane-based; GPS/IMU heading and position estimates support lane keeping and correct intersection turn execution.	Drive 10 known segments; measure heading error at each turn decision point.	Heading error ≤ 5 ; position drift $\leq 1.0\text{ m}$ over segment
Obstacle Detection on Grass	Obstacles appear in open grass with no lane context; pipeline must generalize.	Obstacles at 3 and 5 m on grass; 20 trials per distance.	$\geq 90\%$ at 3 m
Function Completion Rate	Scoring is proportional to functions completed on the Self Drive course.	Functions completed without penalty / total functions per run.	$\geq 80\%$
Zero Operator Interventions	Every estop or reposition is a scoring penalty.	Count interventions per competition run.	0

Section 2: Mechanical Design

The Mechanical Design of Suspended Disbelief is focused on allowing for the robot to have; complete freedom of movement, Easy access to electrical subsystems through doors, robustness, and weather proofing. Suspended Disbelief has custom designed swerve modules, that can strafe in every direction.

The Frame of the Robot was designed to maximize the drive train, and allowing for the smallest robot possible. It's assembled using 80/20 Aluminum, Allowing for the Robot to be lightweight, yet still have strength, and load bearing stability for the Payload. The Acrylic Paneling has built in Hinges to allow for easy access, and the Battery for Suspended Disbelief is housed in the Chassis through a drawer.

Initially, Suspended Disbelief earned its title from the original iteration of the drive system which integrated a four-bar linkage shock suspension system. This season marks our program's first documented attempt at a suspension drivetrain . Due to time constraints, this did not make it to the final iteration of Suspended Disbelief. It remains an experimental initiative with potential for future competitions.

The shock absorber of choice was a 170 mm center distance spring-dampener rated at 750 lbs. This allowed for the spring constant to be adjusted up to 20 mm of axial distance. Regarding the four cross-bars on each module, 1/4"-28 threaded steel rods were cut to 5" and Ball Joint Rod Ends with matching threads. The seated ball provided an ease of assembly with a less restrictive mounting compared to a bearing joint end but more on that later.

As for integration with the rest of the drive system, Suspended Disbelief was equipped with custom designed mounting plates that bolted to the 90° angle brackets at each corner of the swerve modules. Each plate housed four U-brackets for each cross-bar and allowed the rod ends, to reach U-brackets that were mounted directly to the chassis; the choice of threaded rods made for time efficient assembly of the swerve modules to an ideal vertical orientation.

Structurally, the robot could sustain its own weight very comfortably. In a stress test, Suspended Disbelief could reach 150 lbs of static loading before breaching elastic deformation. However, while this was an exceptional success, the degrees of freedom that were allocated through the material selections and linkage design became to grow excessive. Live drive tests demonstrated an unstable pitching motion in the swerve-modules, causing unpredictable moments about the axis normal to the mounting plate. As a result, the wheels of the robot were unable to maintain complete, uninterrupted, contact with drive surface. Following a dynamic analysis, our team concluded that a solution would require a complete redesign of the drive system. As reiteration was deemed beyond our timeline, the suspension project was withdrawn from the competing design.

Despite, Suspended Disbelief lacking suspension, our program has spent the season developing tools to encourage a complete integration in the near future including a mass-spring simulation in MATLAB Simulink for camber angle optimization as well as the custom swerve-mounting plates. The progress made in the suspension project was a milestone in itself and we remain excited to see it soon come to fruition.

Moving on to the final drive system, Suspended Disbelief features four custom-designed swerve modules that are placed in a rectangular configuration on the robot to form the swerve drive base. The modules are a revision of last year's design, with the focus of the revision being an increase of the sturdiness of the modules, and taking the transverse forces off of the drive axle and placing them on a custom designed turret assembly. The modules are made out of 1/4 in. machined aluminum. 3D printed prototypes were made to test the assembly of the modules, and also for

mounting on the robot for software testing while the metal modules were being machined. Each swerve module is capable of turning the wheel 360-degree continuously while providing full driving force.

The modules have a 14.4:1 drive gear ratio for a designed top speed of 5 MPH so that the peak of the motor power curve is at typical IGVC speeds (motors perform optimally at around 60 of their maximum speed). An absolute encoder is installed on the steering shaft and is used to read and maintain the angle of the module between power cycles. Each swerve module features an 6-inch pneumatic wheel, which is ideal for all-terrain and all-weather performance. The modules are designed such that the gearbox is attached to the frame of the robot, so that the modules can be easily removed and installed independently of the gearboxes. This allows for easier assembly and transport, and we can have a spare module fully assembled on standby in case any need to be quickly replaced.




Figure 2.1: A diagram of the swerve module.

2.1 Payload & Weatherproofing

Suspended Disbelief features large acrylic panels for protection from rain, or debris. Weatherstripping is used to provide a tight seal in the 80/20 frame. There are gaskets located on the Paneling

to access motors, and other necessary electrical components outside the frame.

The Payload is specifically integrated at the back of Suspended Disbelief. It has Open Access to the outside, allowing for easy transfer of the Payload whenever necessary.

2.2 Requirement Traceability and Measurements

ID	Design Aspect	Target	Measured	Discussion
M-1: Operating Speed	14.4:1 gear ratio gives ~8 mph mechanical top speed; software caps at 3 mph default with a hard 5 mph limit enforced by the SPARKMAXs.	≤ 5 mph	4.8 mph	Within limit. Dual hardware and software caps mean neither alone can cause an overspeed.
M-2: Ramp Traversal	6-inch pneumatic wheels and rigid direct-mounted modules maintain traction on inclines. IMU-triggered speed reduction applied when climbing.	Capable	15° traversed without stalling across 10 attempts	More than sufficient for IGVC ramp grades.
M-3: Structural Stability	80/20 aluminum frame with direct-mounted swerve modules eliminates the flex from the earlier suspension design.	Visibly stable	No observable camera instability during rough terrain and ramp testing	Rigid frame has held up reliably across all outdoor sessions.

Section 3: Safety

When transporting *Suspended Disbelief* in a vehicle, the 12V lead acid battery is removed and stored separately. The four swerve modules are also detached from the frame during transport to prevent mechanical damage and reduce any risk of movement. Additionally, all latches are locked or zip tied to prevent unnecessary wear and tear.

Outside of vehicle transport, the robot may be driven manually via remote control using Bluetooth. In this mode the vehicle is operated by any team member as long as the estop remote remains in range as otherwise the robot will be unable to drive as the motors will not have power. The battery for the robot is charged outside of the robot using a standard lead acid charger.

The physical EStop button is a clearly marked mushroom head button mounted to the rear of the robot at a height that abides by competition rules. The button is directly wired into our custom EStop PCB which controls a solid state relay.

The wireless EStop system is built around that same EStop PCB which uses LoRa to communicate with the remote. The remote is a small handheld device that must be connected to the robot in order for it to drive. When the remotes estop button is pressed, the transmitter sends the estop command to the PCB via that LoRa radio link. The remote has been tested to have a range of well beyond 300ft, even working in a basement with concrete walls and various turns. Additionally, if the wireless link is lost or no heartbeat is received within a configurable timeout, the EStop PCB is designed to fail safe by triggering a EStop, ensuring that a lost radio link results in a stopped vehicle rather than an uncontrolled one.

Whenever the robot is EStopped, whether that be via the physical button or wireless remote, the receiver PCB cuts power to the dedicated EStop line that enables a solid state relay. All of our motors are wired through the solid state relay such that if the EStop line is ever not driven high (3.3v), the motors will instantly lose power. After an estop, the robot will be unable to drive again until the electronics are physically power cycled which ensures a team member is nearby to ensure all is in order.

Our safety lights consist of RGB LED strips that surround the perimeter of the top of the robot chassis designed to provide 360 degree visibility. The strip is driven by a custom safety lights PCB which receives commands from other devices on the network. Whenever the robot is in autonomous mode, the lights will blink as per competition rules, otherwise they will be solid.

On top of ensuring we meet the requirements of the rules, the safety lights also provide additional debugging functionality. As indicated by Figure XYZ, the color of the RGB strip as well as its current animation mode outside of blinking can be used to determine what state the robot is in. This is incredibly useful to determine if the robot is hitting certain goals we expect it to, like when it reaches a waypoint.

As noted, whenever the robot is being ran autonomously the estop remote must remain in range at all times. If the robot starts to behave in an unexpected fashion, or is on track to cause harm, the estop button is designed to instantly stop the robot. The vehicles maximum speed is limited to the required 5mph in both software and hardware. Whenever the robot is going to enter autonomous mode, not only must it be enabled via the remote but the estop button has an additional purpose which is enabling the mobility of the robot which allows it to actually move.

3.1 Requirement Traceability and Measurements

ID	Design Aspect	Target	Measured	Discussion
S-1: Wireless E-Stop Latency	EStop PCB receives LoRa signal from remote and pulls the shared EStop line low, cutting the solid-state relay instantly.	≤ 1 s (≤ 1 m at 5 mph)	0.3 s / 0.7 m	Well within limit. Fail-safe triggers on heartbeat timeout so a lost link also stops the robot.
S-2: Physical E-Stop Access	Mushroom-head button mounted at the rear at rule-compliant height, wired directly into the EStop PCB.	Rule height; stops robot when pressed	Verified	Clearly visible and accessible from all sides. Confirmed stops robot in all test conditions.
S-3: Safety Light Visibility	RGB LED strips driven by the Safety Lights PCB surround the full chassis perimeter. Blink when autonomous, solid otherwise.	Visibly flashing in sunny conditions	Verified outdoors	Visible at distance in full sun. Also used for state debugging during testing.

Section 4: Electrical and Electronic Design

4.1 System Overview

The team prioritized a modular approach to electronics to keep repairability and reliability high. Suspended Disbelief features several custom PCBs designed in house, all connected through a central Hub PCB that ties into the Jetson Nano via USB-to-CAN. Connections between PCBs are made using 6-pin Molex cables that carry 12 V power, ground, CAN data, and the shared E-Stop signal on a single harness. This keeps wiring clean and makes swapping or debugging individual boards straightforward.

4.2 Computing Platform

Suspended Disbelief runs all onboard computation on a single Jetson Nano running Ubuntu. The Jetson handles the full autonomy stack including camera processing, sensor fusion, and driving logic, and connects to the rest of the electrical system over USB for both the CAN interface and the VN-200 GPS/IMU.

Suspended Disbelief uses three categories of sensors: cameras for perception, a combined GPS/IMU for localization, and wheel encoders for odometry.

SVPRO Global Shutter Cameras We run two front facing cameras used for lane line detection and close obstacle detection in both AutoNav and Self Drive. Global shutter was a deliberate choice as rolling shutter cameras introduce motion blur artifacts when the robot is moving, which causes problems for our vision pipeline. Global shutter eliminates that by capturing the entire frame at the same instant.

VectorNav VN-200 Rugged GPS/IMU The VN-200 is our primary localization sensor for Self Drive waypoint navigation. It combines a GPS receiver and a full IMU in one unit which we feed into a sensor fusion pipeline for accurate position estimates. It connects via USB to the Jetson.

NEO Motor Encoders via SPARKMAX Each of the four drive wheels has a built-in encoder on the NEO motor, read directly by the SPARKMAX controller over CAN. These give us wheel odometry, which we use alongside the VN-200 for localization.

Suspended Disbelief uses 8 NEO motors controlled by 8 SPARKMAX controllers, a component from REV Robotics. The SPARKMAXs communicate over CAN, which runs through our custom CAN Converter PCB. That board translates between our CAN protocol and the REV CAN protocol, so the Jetson can command all 8 controllers from a single USB connection. The controllers are wired in a CAN chain and the CAN Converter PCB sits at the head of that chain.

The vehicle runs off a single 12V lead acid battery that powers both the electronics and the drivetrain. We chose lead acid for its safety, low cost, and easy availability. The capacity gives us far more runtime than the 5 minutes needed for a competition run, so battery life is not a concern.

4.3 Custom PCBs

The CAN Converter PCB bridges our CAN protocol to the REV CAN protocol used by the SPARK-MAX controllers.

The Hub PCB is the central connection point for the electrical system. It connects all PCBs together using a star CAN topology and distributes 12 V power and ground through the Molex harnesses. It also monitors the current draw of each connected device and reports that over CAN, which gives us visibility into power consumption during operation.

The E-Stop PCB manages the shared E-Stop signal that runs across all boards. It listens for the signal from the wireless E-Stop remote as well as the physical button, and drives the E-Stop line accordingly. The line sits at 3.3 V during normal operation and is pulled low when an E-Stop is triggered, which activates a solid state relay that cuts drive power. Every PCB on the robot reads this line, so an E-Stop from any source immediately propagates to the whole system.

The Safety Lights PCB drives the RGB LED strips on the robot from commands received over CAN. This handles the competition requirement for a visible indicator whenever the robot is operating autonomously.

4.4 Requirement Traceability and Measurements

ID	Design Aspect	Target	Measured	Discussion
E-1: Battery Voltage Under Load	Single 12 V 18 Ah AGM lead-acid powers all subsystems. Capacity chosen to far exceed the 5-minute competition runtime.	≥ 11.0 V under full load	11.8 V	0.8 V above threshold with no signs of brownout under peak demand.

Section 5: Perception

Suspended Disbelief uses two forward facing global shutter USB cameras arranged as a stereo pair. They are both calibrated using the OpenCV built in camera calibration utilities (checkerboard based intrinsics and stereo extrinsic calibration) to correct for lens distortion and establish the geometry between the two images. The calibration data is used for two main purposes: in AutoNav, a perspective transform is applied to each camera to produce a top down (birds eye) view of the ground plane in front of the vehicle; in Self Drive, stereo disparity is computed from the image pair to provide estimated metric depth to detected obstacles such as stop signs or barrels.

All perception outputs are tied into a single shared occupancy grid with a cell resolution of $5cm \times 5cm$. Each cell stores an occupancy value that indicates the likelihood an object is located at that position. The grid is roughly centered on the vehicles current position and is updated at a fixed rate.

Occupied cells are inflated by a configurable amount before being passed to any additional logic. This inflation step treats the obstacles as real life sizes so that any path generated on the grid maintains the physical clearance required to avoid obstacles and lane boundaries.

5.1 AutoNav Perception

Lane detection for AutoNav is performed using HSV color thresholding. Each incoming camera frame is converted from BGR to HSV color space and a tuned HSV range is applied to the images to isolate the white lane markings from the pavement. Additional filtering is applied before and after the thresholding to help filter out small gaps and noise.

The binary mask produced by the last step is used to create the top down view using the earlier calibration results. This top down representation can then be used, after the inflation step, for path planning as it corresponds directly to where the lanes are relative to the robot.

Lane detection is determined by using the simple method of determining where the dotted lane is in relation to our two lane cameras. An example of the lane detection, both which lane and where lanes are, can be seen in Figure 5.1.




Figure 5.1: Lane detection determining the lane the robot is in and the type of lane in our simulator.

Pothole detection follows the exact same steps as above, and are projected along side the lanes.

Obstacle detection in AutoNav follows a similar set of steps to the lane detection. A set of HSV thresholds are used to detect barrels, or more specifically detect anything that is not the pavement,

and are then projected onto the top down view. This is used for lane following and basic obstacle avoidance due to its relatively inexpensive cost as compared to running something like a neural network on every frame.

5.2 Self Drive Perception

Lane and painted line detection for Self Drive uses the same HSV thresholding and top down perspective transformation as AutoNav. In Self Drive however, additional data is used along side the top down view to indicate the robot is capable of crossing a dashed lane. This is also seen in Figure 5.1.

As mentioned earlier, our main detection of pedestrians, tires, stop signs, and barrels for Self Drive will be using a YOLO model. Figure 5.2 shows an example of this in action using our simulator.




Figure 5.2: The YOLO model detecting barrels and stop signs in our simulator.

5.3 Requirement Traceability and Measurements

ID	Design Aspect	Target	Measured	Discussion
AN-P1: Lane Line Detection Rate	HSV thresholding isolates white markings; perspective transform projects to top-down view.	$\geq 95\%$ of frames	96.2%	Meets target. Main failure case is shadow transitions; additional HSV tuning has kept misses low.
AN-P2: Obstacle Detection Range	HSV detects non-pavement pixels and projects onto the inflated occupancy grid.	$\geq 90\%$ at 3 m	93%	Meets target. Confidence drops slightly at longer ranges in high-contrast lighting but stays above threshold.
SD-P1: Obstacle Detection on Grass	YOLO detects pedestrians, tires, stop signs, and barrels; stereo disparity provides metric depth.	$\geq 90\%$ at 3 m; $\geq 80\%$ at 5 m	92% / 83%	Meets target. Secondary confidence threshold tuned for grass backgrounds improved reliability.
SD-P2: Lane Classification	Same HSV pipeline as AutoNav with added solid/dashed classification for lane-change logic.	$\geq 95\%$ correct classification	97%	Reliable in both sim and outdoor runs. Edge cases occur when markings are partially shadowed.

Section 6: Driving Logic

The driving logic for Suspended Disbelief is built around a state machine that coordinates all behavior for both AutoNav and Self Drive, although mainly Self Drive. The top level state machine manages simple modes such as determining whether the robot is performing AutoNav or Self Drive tasks. Smaller state machines are created for challenge specific operations.

All driving logic is implemented using the perception described in Section 5. Motor commands are issued through the CAN interface to our CAN converter at a fixed rate. The state machines allow for individual and precise behaviors such as lane following, obstacle avoidance, GPS navigation, etc. Creating these states individually, and integrating a system that allows for states to be mixed together, enables us to test all of the functionality early and in simulation.

GPS position and heading data from the VN-200 IMU/GPS are made available to all states and provide a stable heading reference and position estimate to support lane keeping and intersection navigation.

6.1 AutoNav Driving Logic

Lane following in AutoNav is driven primarily by the top down occupancy grid we produced earlier in the pipeline. When the perception pipeline down the obstacle inflation, it essentially allows us to do lane following directly as the markings are relative to the vehicle with real distances. When both lanes are visible the driving logic computes the midpoint between the left and right lane boundaries at some fixed lookahead distance in front of the robot, and generates a point combined with any GPS influence. We then use Peer Pursuit to create a path that the robot will follow for a fixed time before updating it again. There are precautions such that if a lane or obstacle is too close to the robot, it will override the current path in an effort to avoid it. When only one lane boundary is visible the driving logic falls back to just maintaining an offset from the visible lane.

$$\begin{aligned} Goal_x &= center + (goal_{offset} \times goal_{weight} \div goal_{distance}) \\ Goal_y &= robot_{center} + offset \end{aligned}$$

Obstacle avoidance is handled by an A* planner operating directly on the inflated occupancy grid. When the perception pipeline marks cells as occupied by a detected obstacle, the inflated footprint is treated as impassable and thus the robot will search for the shortest path around all obstacles. The goal for A* is at a fixed distance ahead of the center of the robot, however, it is also slightly influenced by the distance to the next waypoint. The algorithm that determines the horizontal point of the goal can be seen below.

The ramp is handled implicitly through the vehicles physical configuration and lane following abilities. In software the vehicle does not require any additional ramp detection so the lane following and GPS navigation continue as intended even on the ramp. When the IMU detected the robot is driving up an incline, a small decrease in speed is applied so that the robot safely crosses the peak and also reduces camera shake.

As noted above, GPS waypoint navigation runs continuously along side lane following and obstacle detection. Our VN-200 provides real time position estimates which are compared against the current target waypoint if applicable. The goal selection as described in the Lane Following section explains how the required heading to the next waypoint impacts the goal position, along

side the distance from the waypoint so the robot is less inclined to drive towards it at farther distances.

Vehicle speed is controlled by multiple locations; the SparkMax controllers, the physical swerve modules, and the software. At the software level, the driving logic applies a speed reduction based on a few core conditions: the pitch of the robot, the distance to the nearest obstacle, and when any high curvature planned path is generated. Each of these can compound to limit the robots maximum speed, although to be safe a minimum speed is also set to prevent all three from allowing the robot to move. The default maximum speed is set at 3mph, although, it is easily tuned through configuration.

The output of all of the driving is a desired forward velocity, sideways velocity, and angular velocity. The forward and sideways velocity describe how the robot is going to translate, and the angular velocity describes how the robot is going to spin. The default behavior is for the robot to aim itself at the goal point and maintain a forward velocity, the sideways velocity when driving autonomously is not currently used. Each modules output is regulated using a PID onboard the SparkMax motor controllers, which are instructed by the CAN converter PCB.

6.2 Self Drive Driving Logic

Lane keeping in Self Drive uses the same center line tracking system as AutoNav. The Self Drive state machine differs however as it has the ability to track which lane the robot is currently in so that it knows what direction it is able to change lanes (as determined by the solid vs dashed lines).

Lane changes are triggered by the state machine whenever a obstacle is detected in the current lane. In this case, the state machine will also confirm the adjacent lane is unoccupied and triggers a lane change, which marks the dashed lane as not an obstacle and nudges the A* goal point in that direction. If both lanes are determined to be invalid paths, the robot will back up and try again. Whenever a lane is fully changed, determined by the distance between the robot and the dashed lane, then state will end and the robot will continue normal behavior.

When the state machine determines an intersection is present (typically through a stop sign and knowledge of the current task), the driving logic will initially make a stop at the stop sign if applicable. The small white horizontal lane near the stop sign directly in front of the robot will be marked as a non obstacle.

After stopping at the stop sign for a specific amount of time, the robot begins driving through the intersection and triggers a turn, if applicable. If a turn occurs, the robot will prioritize the lane on the right hand side if detectable, otherwise whichever lane it is capable of entering.

Compared to other obstacles, when a pedestrian is detected by the YOLO pipeline within some configurable stopping distance, the robot will come to a stop and remain stationary until the pedestrian is no longer detected in the front of the robot. Depending on the state of the state machine, and knowledge of the distance of the pedestrian, the robot may also choose to change lanes if possible.

All three parking types are implemented as scripted sequences triggered by the state machine when the corresponding task is active.

Pull out and **pull in** maneuvers use a combination of lateral swerve translation and forward/reverse motion. Because the swerve drive can strafe without rotating, the vehicle can pull directly into or out of a parking space perpendicular to the road with minimal maneuvering.

Parallel parking uses a preplanned motion sequence: the vehicle first aligns with the target space using lane context, then performs a reverse arc. The absolute encoders on each swerve module provide real time steering angle feedback to ensure the sequence executes as planned.

Potholes are detected by the perception pipeline as circular obstacles projected onto the occupancy grid (Section 5). When a pothole is detected in the current lane, the Self Drive state machine evaluates whether a lane change is available. If one is, the lane change behavior described earlier is triggered and the vehicle changes lanes to avoid the pothole. Once past the pothole location, the vehicle returns to the original lane. If no lane change is available, the vehicle reduces speed and if possible steers around the pothole within the current lane using the A* planner.

Stop sign detections from the YOLO model are subject to a secondary OCR validation step before the driving logic acts on them. After YOLO identifies a stop sign bounding box for a configurable amount of time, the cropped region is passed to an OCR engine which attempts to read the text content of the sign. A detection is only treated as a valid stop sign if the OCR output matches the expected “STOP” text with sufficient confidence. The color range of a stop sign is also taken into account to validate it is a standard red stop sign.

6.3 Requirement Traceability and Measurements

ID	Design Aspect	Target	Measured	Discussion
AN-D1: Lane-Center Tracking Error	Lane-following computes the midpoint between left and right lane boundaries at a fixed lookahead; Pure Pursuit generates the path at 10 Hz.	MAE \leq 15 cm; worst-case \leq 30 cm	MAE 11.3 cm; worst-case 24 cm	Both within target. Worst-case excursions occur at the sharpest test course turns.
AN-D2: Obstacle Avoidance Clearance	A* planner on the inflated occupancy grid treats obstacle footprints as impassable. Inflation radius tuned to vehicle half-width plus safety margin.	\geq 15 cm clearance; zero contact events	18 cm min; zero contacts	Consistently above clearance target across all configurations.
SD-D1: Stop Sign Stopping Precision	Intersection approach state decelerates to a configurable stop distance from the YOLO-detected sign, confirmed by OCR validation before acting.	Stops within 30 cm of stop bar	22 cm average	Within target. OCR validation prevents false stops from non-standard signs.
SD-D2: Fake Sign Rejection	Two-stage pipeline: YOLO bounding box held for configurable duration, then OCR must read “STOP” above confidence threshold before intersection behavior triggers.	Zero false stops per run	Zero across all test runs	No false positives observed. Color validation provides a third layer of confirmation.

Section 7: Key Performance Indicators

7.1 Overview

We selected KPIs that map directly to competition scoring outcomes. For AutoNav, that means completing laps cleanly within the time window without violations or obstacle contact. For Self Drive, it means executing tasks tasks successfully.

7.2 AutoNav Key Performance Indicators

Table 7.1: AutoNav KPI Summary.

ID	KPI	Target	Measured	Units
AN-K1	Penalty-free course completions	≥ 2	2	count
AN-K2	Mean lap time	≤ 4	3.4	min
AN-K3	Lane line detection rate	≥ 95	96.2	%
AN-K4	Obstacle contact events per run	0	0	count
AN-K5	Lane-center tracking MAE	≤ 15	11.3	cm

7.3 Self Drive Key Performance Indicators

ID	Measurement Method	Target	Measured	Discussion
SD-K1: Function Completion	Functions completed without penalty divided by total functions attempted per run.	$\geq 80\%$	85%	Above target. Missed functions traced to GPS heading drift causing late turn initiation at one corner of the test course.
SD-K2: Operator Interventions	Count estop presses and manual repositioning events per run.	0	0	Zero across all test runs. Clearest indicator the autonomy stack is competition-ready.
SD-K3: GPS/IMU Heading Accuracy	Drive 10 known course segments; record heading error at each turn decision point and position drift at segment end.	Heading ≤ 5 ; drift ≤ 1.0 m	3.2 avg; 0.74 m	Within target. Remaining drift is primarily GPS multipath near the tree line.
SD-K4: Obstacle Detection	Place obstacles on grass at 3 and 5 m; run detection pipeline 20 trials per distance.	$\geq 90\%$ at 3 m	92% / 83%	Meets target. Secondary threshold profile tuned for high-contrast midday lighting.

Section 8: Analysis of the Complete Vehicle

8.1 Lessons Learned During Construction and Integration

The biggest lesson was to validate mechanical concepts before full integration. The four-bar suspension consumed significant machining budget before dynamic testing revealed the ball joint rod ends introduced uncontrollable module tilt. We direct-mounted the modules to the chassis as a corrective action and recovered reliable drivetrain performance. The suspension work carries forward into next year's design.

8.2 Software Testing, Bug Tracking, and Version Control

All software is managed in a GitHub repository using a branch and pull request workflow; changes merge only after peer review. Bugs found in simulation or physical testing are logged in GitHub Issues and linked to the resolving PR. Perception is validated by reviewing occupancy grids and detection overlays on recorded runs; driving logic is validated by observing full autonomy runs in simulation and on the physical test course.

8.3 Simulation Based Testing

We built a custom simulator that runs the full autonomy stack against simulated hardware — CANbus, PCBs, VN-200, and cameras. It supports configurable course layouts, lighting, and obstacle placement, letting us iterate on perception tuning and driving behavior without needing the physical course. Scripted repeatable scenarios (pedestrian mid-run, fake stop sign alongside a real one) have been especially useful for edge case testing.

8.4 Physical Testing

Physical testing takes place on a parking lot course with taped lane lines and cones approximating the AutoNav layout. Behaviors validated in simulation have transferred to the physical robot with minimal changes — mainly HSV threshold adjustments for real-world lighting. This close sim-to-real performance gives us confidence heading into competition.

8.5 Overview

Suspended Disbelief's main attack surfaces are its exposed Ethernet port, unauthenticated CANbus, and third-party model inputs. None of the onboard communication protocols implement authentication or encryption, which would need to be addressed before any production deployment.

8.6 Vulnerability Summary

Table 8.1: Cyber vulnerability summary.

#	Vulnerability	Hardening Steps
1	Exposed Ethernet port, accessible through any open panel with no access control. An attacker with brief physical access could modify software or configuration.	Physical port lock during operation; default-deny firewall; key-only SSH.
2	Unauthenticated CANbus, any device attached to the bus can send arbitrary commands to motors or PCBs.	Physical access controls; message signing and whitelist filtering on safety-critical nodes.
3	Sensor and model inputs, YOLO and OCR pipelines could be fooled by adversarial physical alterations; ONNX model is a supply chain risk.	Additional training data for robustness; hash verification of loaded model files.