Camera

- **What is it?**
  - Video cameras can be mounted at certain parts of vehicles to aid the driver or the system in the driving task
- **How it works**
  - ***Show or Draw Slide 1: Camera locations on vehicle and FOVs***
  - Cameras can have varying focal lengths depending on the need
    - Fish eye cameras (wide angle lenses) have wide FOVs, often almost 180 degrees
      - Good for seeing things that are close to the vehicle
    - Cameras with narrower FOVs than wide angle can give video feeds of objects further away
  - ***Shows Slide 2: Camera positions on Tesla***
  - Camera resolution has risen over the years to provide high resolution and high definition feeds if needed
    - Can use 1080p or better camera with 8 megapixels or more, typical camera has at least 2 MP
    - Higher resolutions need greater bandwidth, need capable systems and GPUs
  - ***Show Slide 3: Some examples of automotive cameras, both from Conti***
  - Camera feeds video data over serial port or similar interface
    - Usually capture video in the RGB color space for each pixel
  - Distances can be determined and objects can be detected using some camera set ups
  - Many cameras feed directly into the infotainment system of the vehicle
  - Very cheap to manufacture camera parts
  - Machine vision and machine learning is key for camera vision
    - Tesla has notoriously avoided lidars and focused on using camera vision instead
  - Cameras can be susceptible to glare, bright sunlight, and weather conditions though
- **What are some camera configurations and uses?**
  - Rear camera
    - ***Show Slide 4: Live view of Tesla rear camera on front stack***
    - Covers blind spot when reversing
    - Rear vehicle collision warning
    - Rear cross traffic warning
    - Rear pedestrian alert
    - Trailer hitch assist
  - Forward facing cameras
    - Lane line identification
    - Edge of road detection
    - Traffic sign identification

        - Lane departure warning
        - Lane keep assist
        - Rain sensing
        - Forward collision warning
        - Pedestrian detection
        - AEB
        - Landmark detection
        - ***Show Slide 5: AEB with camera***
        - ***Show Slide 6: Video from conti of lane departure warning***
    - Surround view cameras
        - Placed along the sides or corners of vehicles and have wide angle lenses to cover as much a FOV close to the car as possible
        - Blind spot detection
        - Lane change assist
        - Pedestrian detection
        - Parking spot detection
        - Birds eye view
        - Mirror replacement
    - In-cabin cameras
        - Driver monitoring systems
            - Determine if driver is wearing seat belt
            - Looking away from the road
            - Adjust airbag deployment
            - ***Show Slide 7: Head pose estimation gifs***
        - Occupant monitoring systems
            - Determine if children and pets are present
            - Rear seat child unattended warning
- **What signal processing is needed?**
    - Cameras need to be carefully calibrated when installed on vehicles, especially for ADS
    - The video frames are passed to powerful compute systems, such as GPUs, especially when multiple cameras are used
        - This system needs to transform the video frame based on the FOV of the camera (fish eye effect)
        - ***Show Slide 8: Wide angle lens distortions and fixes***
        - Transform the video into other color spaces to increase contrast and aid in lane line detection if necessary (Hough transformation)
            - Can also be used for parking line detection
            - Can even project lane lines forward if the line of sight is blocked
            - ***Show Slide 9: Hough Transformations***
        - Implement machine learning algorithms to identify objects such as other vehicles, trucks, pedestrians, bicyclists, motorists, and the road itself

            - Need good models of these objects to train the neural nets on to accurately identify objects while driving
            - OpenCV and other machine learning frameworks leveraged for this
        - ***Show slide 10: One frame of some roadside object identifications***
        - Use machine learning and calibration to determine distance each object is away from the vehicle
        - ***Show Slide 11: Tesla Autopilot camera vision video**

GPS

- **What is it?**
    - ***Ask what gps stands for, or what GNSS stands for, and how it works***
    - Global Positioning System, also referred to as GNSS, global navigation satellite systems
    - ***Show Slide 1: Picture of Novatel and generic GPS antenna***
- **How does it work?**
    - Orbiting satellite constellations can triangulate your position on the earth using transmitted signals
    - ***Ask if anybody knows anything about GPS satellites***
    - ***Show Slide 2: GPS satellites and orbitals***
        - Requires a GPS antenna/receiver to convert signals and calculate positions
        - Usually 4 satellites need to be visible overhead for GPS receiver to calculate a 3D position on earth
        - GPS signals are actually microwaves in the "L1" or "L2" EM band, 1575 MHz and 1227 MHz respectively
            - FM radio is between 87.5 and 108.0 MHz and Wi-Fi operates at 2.4 or 5.0 GHz
            - GPS signals have wavelengths of around 15-30cm that can travel through parts of the atmosphere better than some other frequencies
    - The GPS satellites used carry precise atomic clocks and know their own position over the earth while orbiting
        - There are 27 satellites spread out over the sky for the "core" satellite constellations of GPS maintained by the DoD, usually 6-8 are capable of being seen overhead at any given time
        - Maintain an orbiting radius of about 26,600 km or 16,500 miles from the center of the earth, rotate around the earth twice a day and cross over approximately the same spots at the same times each day
            - They move at over 14,000 km/h
            - Interesting fact, the clocks on GPS satellites actually tick about 38 microseconds faster than those on earth due to special relativity

(less gravity at high orbits), the satellite's clock take this into account

- ***Ask anybody if they know how GPS receivers estimate position***
- Essentially, GPS receivers use the time-of-flight rule to identify how far away it is from each satellite it is connected to
    - ***Show or draw Slide 3: time-of-flight with gps satellites***
    - For a simplified explanation, 3 satellites can give 2D position, add a 4th and you can get elevation measurements
    - ***Show or draw Slide 4: Satellite triangulation/multilateration***
    - Satellites emit a carrier wave for a GPS receiver to connect to
    - It also adds some extra modulations that let the receiver determine at what time exactly the signal was transmitted from the satellite so an accurate time of arrival can be deduced
    - Positions are commonly reported in latitude, longitude, and elevation, usually height above mean sea level
        - Can convert to local vehicle frame if needed
    - Positions can then be tracked over time for greater accuracy by the receiver
        - Velocity estimates are available as a result of the tracking over time
- Receivers automatically keep track of the satellite constellations above
- Does not require an internet connection
- GPS signals are highly standardized by governing organizations (NMEA specification)
- ***Ask if they know any other countries GNSS satellite names***
- Other nations have their own GPS satellites associated with GNSS that a receiver can connect to as well
    - Beidou (China)
    - GLONASS (Russia)
    - Galileo (European Union)
    - QZSS (Japan)
    - NavIC (India)

- **How accurate is GPS?**
    - Using basic standards, a GPS 3D coordinate is usually accurate to within 15m (49ft) with an unobstructed FOV to the satellites being used
        - Can be slightly better during tracking, but there is always at least 5m of error, if not more
- **What are the sources of error?**
    - ***Ask if they can think of any sources of error***
    - Mostly due to EM waves distorting as they travel through the earth's atmosphere
    - ***Show or draw Slide 5: atmospheric signal interference***
        - Weather conditions can introduce random "noise" to the signal or degrade it overall

        - The top layer of the atmosphere (the ionosphere) starts around 250 miles above the surface of the earth
        - Signal is affected by changing its path, slowing it down, and altering the waves phase
        - Introduces the 5-15m error into the signal
    - Poor line of sight can result in bad signal strength
        - ***Show or draw Slide 6: Buildings blocking signal***
        - Mountains, buildings block a lot of the signal, usually can't use a gps receiver indoors
    - Metal objects, electronic components or other objects of high density can block or reflect signals
- **What can we do to minimize positional errors?**
    - Ensure unobstructed view to the sky
    - Ensure up to date receivers
    - Use differential GPS (DGPS)
    - ***Ask if anybody knows what DGPS is or stands for***
- **What is DGPS?**
    - ***Show or draw Slide 7: DGPS process overview***
    - DGPS uses the fact that the atmosphere introduces random noise to the GPS signal to its advantage by using two GPS receivers (a base and a rover) within close distance (< 100 miles usually) that can talk to each other
        - The base is set up as a reference station (not moving)
            - In initial set up the base conducts a multi hour survey (usually at least 24 hours)
            - As random noise comes in to the base over the GPS signal while surveying, the base station can actually start to see where its actual position is on the earth with high accuracy (< 20cm)
            - Remember the average value of truly random noise should average to 0 over a long enough time period, dgps exploits this
            - ***Show or draw Slide 8: GPS vs DGPS accuracies and survey time***
        - After this set up, the base station knows its actual position on earth and can calculate any errors in position that the GPS satellites are sending in real time
            - ***Go back to Slide 7 for DGPS process overview***
            - The base then passes over the "differences" in received signal position vs its known actual position to the rover GPS receiver, which is connected to GPS satellites as it normally would
                - These are the "corrections", sent usually as radio waves
            - The rover then uses the corrections to fix the position it receives from the GPS satellites ultimately leading to higher accuracy, sometimes all the way to 3cm

            - Lower end systems can still get accuracy less than 1m if not better, many can get errors under 20cm
        - This works because the atmosphere/weather above the two gps receivers is mostly the same as long as they are within 100 miles, so the corrections are valid for both receivers
    - The DGPS base station continuously operates, broadcasting differential corrections for any GPS receivers in the area to connect to (if they have the capability)
    - Specialty software and equipment are needed, there are many high-end systems
    - Can operate at 20Hz or higher
- **What is GPS used for vehicles?**
    - ***Ask if anybody knows what GPS/DGPS is used for***
    - Localization within maps
        - Determine when certain L2 features can be used
    - Navigation
        - Instructions for the driver
        - Global route planning for automated driving systems
    - Speed estimation
    - Traffic congestion estimation

IMU

- **What is it?**
    - ***Ask if anybody knows what it stands for and what an IMU does***
    - Inertial measurement unit
    - ***Show Slide 1: IMU examples from adafruit, Honeywell, vectornav***
- **What does it do?**
    - Measures acceleration, angular speed, and orientation of whatever it is mounted on
    - ***Ask if anybody knows how it works***
- **How does it work?**
    - IMUs use accelerometers, gyroscopes, and magnetometers micro machined onto silicon wafers to measure the kinematic and dynamic state of the vehicle
    - ***Show Slide 2: IMU Components***
    - Makes use of MEMS (micro Electrical Mechanical System) technology
    - Usually have 9 DOFs
        - ***Ask if anybody can say what the 9 DOFs are and how they relate to a vehicle***
        - ***Show Slide 3: Vehicle frame***
        - Accelerometers can measure acceleration in the x,y,z directions
            - Make use of really small springs that change capacitance when subjected to accelerations
            - Output acceleration values, m/s^2
            - Acceleration values can be converted over time to velocity in each direction

        - Gyroscopes measure angular rotation rate around the x,y,z axes
            - ***Show Slide 4: Gyroscopic data gif and graph***
            - Output measurements in degrees or radians per second
            - Rotation rates can be tracked over time to determine angular position around the axes
        - Magnetometers detect heading angles for the three axes
            - Output is in degrees or radians
            - Use microscopic coils to detect changes in magnetic field as sensor is rotated
    - Can be updated at 1,000 Hz or more
    - High end systems have extremely high resolutions for all three measurement systems making them highly accurate
        - All IMUs "drift" over time though but can be managed by combining other data sources
    - The values from the accelerometers, gyros, and magnetometers are usually passed to a central system (sometimes referred to as an Inertial Navigational System, INS) that ultimately tracks
        - Velocity in the x,y,z directions
        - Orientation (i.e. roll, pitch, yaw) and heading
        - Will most likely use quaternions to avoid gimbal lock when calculating positions and angles based off rotation data
    - High end systems can actually perform dead reckoning navigation based only IMU data
        - ***Show or draw Slide 5: Dead reckoning overview when GPS drops out***
        - Uses kinematic and dynamic equations to track where the vehicle is and how fast its going based off a starting position and speed and the accelerometers, gyros, and magnetometers
        - Can often track a vehicle for up to 10 seconds with less than 30cm errors
    - IMU readings and the INS output are usually combined with other navigational systems, like the GPS, for greater accuracy
        - Dead reckoning can help when GPS drops out or has a poor signal
    - IMU is not affected by wheel slip or skid and can still give the velocity and direction during these events
    - Almost all vehicles have some sort of IMU on board as they can play integral part in stability control
        - May only focus on yaw, and accel. In x and y
        - Yaw rate (rotating around the vertical z axis) is important to know for different dynamic controls
        - Usually placed somewhere near the middle of the vehicle, under center console between driver and front passenger seat
    - IMUs have a number of advantages
        - Cheap
        - Simple to set up

        - Reliable
        - Is environment independent (does not need to interact with external environment like radar, camera, or lidar does. This leads to fewer failure modes)
- **What is IMU used for?**
    - ***Ask if anybody knows what IMU is used for in vehicles***
    - Vehicle dynamic control
        - ***Show or draw Slide 6: Oversteer understeer, IMU can detect actual velocity of the vehicle***
        - Electronic stability control
            - If a vehicle is understeering or oversteering during a turn the vehicle CPU can compare wheel speed, steering angle, and throttle position to the readings of the IMU to determine the vehicle's actual movement
            - Vehicle determines which wheels need to be braked to put the car back on its target line
                - Or throttle back engine
        - Rollover detection
    - Crash sensors use accelerometers to detect collisions
    - Dead reckoning position and velocity estimations
        - Can be used on its own, or combined with other sensors for greater accuracy
        - Aids with localization within HD maps
    - Aids with automatic lane keeping and other automated driving tasks
    - ***Show or draw Slide 7: IMU accel data from a crash***
    - Crash reconstruction or data analysis of crash or near crash events

LiDAR

- **What does it stand for?**
    - Light Detection and Ranging
    - ***Show Slide 1: pictures of lidars from different manufacturers***
- **How does it work?**
    - *** ask if anybody know how it works***
    - ***show Slide 2: EM spectrum with LiDAR cutout***
    - Much like radar, lidars use pulsed lasers that reflect off objects to perceive the environment around us
    - Lasers are pulsed at thousands or millions of times per second and the time of flight principle is used to figure out ranging and position, all in real time
    - Wavelengths used in automotive lidars are 905 nm and 1550 nm (nano = 10e-9)
        - This puts the light used in the "near infra-red" category on the EM spectrum
        - Most operate around 900 nm, safer against eye damage
    - Each laser can detect objects in its "plane" based on the orientation (you can have the plane pointed up, pointed down, or flat

- To maximize FOV, manufacturers include multiple lasers positioned at different angles to cover areas such as the ground, what's directly ahead/behind you, and anything overhead
  - Typical vertical FOV is +/- 30 degrees
  - Multiple lasers are referred to as channels or scans
  - Use only one channel and you have a laser scanner
- *** Show or draw Slide 3: Vertical channel separation***
- *** Show or draw Slide 4: Channel separation comparison***
- Add multiple channels to a revolving sensor and you can achieve a 360 degree horizontal FOV that covers +/- 30 degrees vertically
  - In modern lidars each channel can detect up to about 4,000 points per revolution
    - Often at speeds of 10-20 Hz
  - Modern lidars also have different numbers of channels
    - 16, 32, 64, or 128
    - More channels mean more resolution and more coverage
    - The farther away from the sensor the more "spread" out each channel's plane gets, many channels focused around +/- 5 degrees
    - The number of returned points approaches 10 million per second if you have a 128 channel lidar
- Some lidars do not spin but have a fixed horizontal FOV between 30-150 degrees with multiple channels
- Lidars usually have optical filters in front of the sensor (such as a lens) that filters out sunlight or other wavelengths close to the one emitted from the sensor
- Mechanical lidar is the traditional lidar (rotating/scanning)
  - New advanced technologies are emerging
    - Flash lidar
    - MEMS (quasi solid state) that bends the lasers without moving optical components
    - Solid state/digital: uses thousands or millions of small lasers projected onto the scene without the need of laser bending

- **What type of data do we get from lidars?**
  - *** ask if anybody has seen lidar data visualized***
  - *** Show Slides 5&6: point cloud examples***
  - Each pulsed laser can return a point with an x,y,z and intensity value (4 dimensional measurement). Here the intensity refers to the power intensity of the returned signal, represented by change of color when visualized
  - Each point can be added to a "point cloud" for the frame
    - When visualized, ***outlines*** of objects are often visible
      - Cars, trucks, other vehicles
      - The road way, lane markings, adjacent terrain
      - Road signs, barriers, guard rails, traffic lights
      - Fences, trees, pedestrians

            - ***Show Slide 7: Shape outlines***
            - Intensity values vary based on the material it bounces off of and the angle it hits the object at
                - Rubber doesn't reflect well, water can cause problems
                - Flat metallic objects or other hard materials reflect light well when facing the sensor
                    - But when orientated at an angle, intensity of the reflection can be lower
                - Clothing, organic material, have lower reflectivity but can still be detected consistently
- **What are the lidar data specifications?**
    - For a 32 channel Velodyne Ultra Puck LiDAR
        - Range: 200m
            - Range accuracy: +/- 3cm
        - Horizontal FOV: 360 degrees
            - Horizontal angle resolution: 0.1 degrees
        - Vertical FOV: -25 – 15 degrees (40 degrees total)
            - Vertical angle resolution: 0.4 degrees
        - Frame rate 5-20Hz
        - Number of points per second: 600,000 – 1,200,000
        - Power consumption: 10W
        - Weight: 925 grams
- **What are the processing requirements to use lidar data?**
    - ***Slide 8: Some lidar electronic component drawings***
    - Typical specialty electronics and CPUs are used in lidars
        - Laser driver
        - Motor to spin the sensor
        - Electronic filters
        - Analog to digital converters
        - Digital signal processing
        - Power management
        - Interfacing to vehicle
    - Data is natively output in point clouds. Each point has 3D position and then intensity value
        - Visualization is only for humans
    - Due to large number of data points, down sampling almost always occurs
        - Can use voxels, octrees, height maps or other data structures with slightly lower resolutions
    - Identify ground plane
        - Can use a simple height cut off if on level ground
        - Or a popular plane estimation algorithm such as RANSAC (random sample and consensus) to identify the ground
            - Can identify sloped roads better

        - Ground can be removed or stored in a separate point cloud
    - Cluster the point cloud for certain objects
        - Can be based on Euclidean distance or density values of the points in space
        - In clustering we want to identify all the points that come from important roadway users or objects such as
            - Other vehicles
            - Pedestrians and cyclists
            - Guard rails, fences, curbs
        - These objects will usually be picked up by multiple channels of the lidar
        - Specialized algorithms and data structures help with fast proximity searches for clustering
        - Cluster points can be identified, centroids calculated of the object, and bounding boxes can be drawn that better resemble the 3D shape of the object in the point cloud
        - *** Show Slide 9: Bounding boxes in one frame of a point cloud***
        - Object classification can be determined sometimes as well from the clusters
    - Track clusters over time using Kalman filters
        - Based on the extracted clusters for each frame, the objects can be identified and tracked over time using the bounding box, centroid, and object classification if available
            - This helps the vehicle understand better where objects are around them, where the drive-able area is, and where objects around them are headed in the near future
        - It's common to track the velocity, orientation, and angular velocity of these clusters/objects
    - Localization
        - Other objects, such as trees, buildings, traffic signs, and other landmarks can help localize vehicles or build maps or certain areas
            - Other disciplines use lidar technology for mapping as well
            - Civil engineering and construction companies can use lidar to map work sites
            - Archeologists/environmentalists/geologists map whole areas of the land
- **What ADAS features do lidars help with?**
    - *** ask if anybody knows what LIDAR could be used for in the driving task***
    - Still a "future" technology and not implemented on any production vehicles today
    - They play an integral part in automated driving systems (ADS). Lidars can inform
        - Object detection including heights
        - Object classification
        - Object tracking

        - Lane marking identification
        - Localization and mapping
    - These would then feed into automated driving algorithms that perform the driving task, such as
        - Braking, accelerating
        - Lane changes, lane following
        - Object avoidance
- *** Show Slide 10: Velodyne lidar video to sum up the lidar experience***

Localization techniques

- **General Navigation Maps**
    - Typical navigational maps that your GPS system may use
        - ***Show Slide 1: normal navigational maps***
        - Google maps
        - Apple maps
    - Contains road way information, city/country/state information, location of stop signs, traffic lights, speed limits, business location information, and may integrate traffic in real time
    - With a normal GPS signal can locate yourself within 10m or so given a good view of the sky
    - May be good for general navigation, but as automated driving systems capabilities increase, we need better maps with more detail
- **HD maps**
    - ***ask what other information might be useful from a map for the vehicle***
    - HD maps include much more detailed information focused on localization
    - Include DGPS, lidar point clouds, and camera data of vehicles that have already "mapped" the area
    - Consist of dense point cloud data as well as layered roadway infrastructure and lane boundary locations
        - Locations of lane boundaries
        - Locations of road boundaries
        - 3D locations of road signs
        - 3D locations of traffic lights
        - Cross walk locations
        - Road curvatures
        - Roadway segments such as the roadway itself, the roadside, barriers, nearby foliage and trees
        - ***Show Slide 2 and 3: HD map pictures***
    - A vehicle can use these maps to compare what its sensors see to what is expected from the maps at its location
        - The vehicle can maintain lawful driving and position itself in the center of the lane at all times
        - You can match up point cloud data to estimate position using localization and mapping techniques

        - You can also estimate the distance to landmarks using camera and lidar
            - Triangulate position based on distance from landmarks marked in the maps along with positions from the DGPS and IMU systems
    - HD maps are essential for automated driving systems localization and path planning algorithms
    - However, they are extremely hard to make and store
        - Proprietary data
        - Time consuming to collect data
        - Large amount of data
        - Storage or transmission to the vehicle to use is challenging
        - HD maps have to be updated as the roadway changes (which it always is all over the country)

Other Sensors

Wheel Speed Sensor

- **What is it and what does it do?**
    - Sensor that measures wheel speed at the tires
    - ***Ask if anybody knows how they measure wheel speed***
- **How does it work?**
    - Usually based off Hall effect sensors
        - ***ask if anybody knows what Hall effect sensors are***
        - Utilize changing magnetic fields to produce a signal
        - ***Show or draw Slide 1: Hall effect sensor arrangements for wheel speed***
    - A magnet wrapped in a wire coil is placed near the axle of the wheel and the sensor
        - A toothed ring or a multipole ring is installed to rotate with the axles as the wheel moves
    - As the toothed or multipole ring rotates near the magnet and coil, the magnetic field will change in a quantifiable way based on how fast the teeth or poles are passing the sensor
    - ***Show Slide 2: Some examples of wheel speed sensors***
    - Microcontrollers connected to the sensors can convert the electrical “ticks” into wheel speed

Anti-lock braking system

- **What is it?**
    - ABS, or anti-lock braking system
- **What does it do?**
    - stops the wheels from skidding in a high braking event
- **How does it work?**
    - ***Ask if anybody knows how ABS works***

    - Based on wheel speed sensors and sometimes IMU data, the ECU can modify braking pressure to all four wheels if a lock-up is detected
    - ***Show or draw Slide 3: ABS system components and pulsing brake force graph***
    - Piggy backs off the wheel speed sensor (ABS and wheel speed sensors are often integrated together)
    - Uses the hydraulic brake pump system to quickly modulate brake pressure when skidding is detected, up to around 30 times per second
    - Less skidding means more traction at your wheels and quicker stopping times while reducing collisions, increasing control, and saving lives

Wheel Angle/Steering Sensor

- **What is it?**
    - Wheel angle and steering angle sensors
- **What does it do?**
    - Measures the steering angle, a.k.a. the desired steering angle from the driver
    - Measures the actual wheel angle produced by the steering system
- **How does it work?**
    - Multiple sensors located in the steering column and the power steering system
    - ***Show Slide 4: Steering angle sensor components and diagram***
    - Most steering angle sensors these days are digital
        - Some version of a shaft encoder that sends steering angle position and how fast the steering wheel is being turned to the ESC
        - Steering torque can be measured as well (how hard the driver is torqueing the steering wheel)
    - Power steering system uses similar sensors to report the wheel angle at any given time
        - Most power steering systems are still mechanically connected to the steering column in some way, there are a few instances of drive by wire

ESC, Electronic Stability Control

- **What is it?**
    - Electronic stability control
- **What does it do?**
    - Modulates power or braking pressure to individual wheels during sharp turns and/or braking events
- **How does it work?**
    - ***Ask if anybody knows how the ESC system works***
    - Uses the wheel speed data, steering angle data, wheel angle data, and data from the IMU to detect when the vehicle has lost control
        - ***Show Slide 5: ESC system overview***
        - During understeer and oversteer, the wheel angles and vehicle velocity do not match up with the intended steering angle from the driver
        - Tires may be losing grip with the road from skidding or slipping

- The ECU can modulate brake pressure to the individual wheels and engine power to try and return the vehicle to the intended driving line
- Most likely needed in high speed turning events such as swerves that may also be accompanied by hard braking from the driver
- Mandated in all vehicles from 2012 and after, though many vehicles already had them installed by this time
- Especially important for SUVs
- Estimated that vehicles with ESC are involved in around 30% fewer single vehicle crashes and around 60% fewer rollover crashes
- ***Show video on Slide 5: euro NCAP ESC test from 2012***

Throttle Position Sensor

- **What is it and what does it do?**
  - Sensor that detects how far down the throttle (i.e. gas pedal) has been pressed
- **How does it work?**
  - Older vehicles used mechanical systems that used potentiometers to determine how far down the driver has pressed the gas pedal
    - As the pedal moved down the electrical resistance in the signal decreased in a quantifiable way
    - The ECU could then determine how far the gas pedal was depressed from the potentiometer signal
  - New vehicles use electronic throttle control (ETC) and accelerator pedal position sensors (APP) to determine throttle position
  - ***Show or draw Slide 6: ETC vs mechanical throttle position sensors***
  - ETC and APP are digital and replace the more analog system of steel cables and potentiometers
  - ETC and APP use the Hall effect sensors to determine position
    - Covert changing magnetic fields to voltages within the sensors
    - Pass these signals to the ECU to determine how far the accelerator has been pressed and how much to open the throttle valve
  - ETC systems allow for the easy integration of cruise control, adaptive cruise control, AEB, and ESC as it is all controlled digitally by the ECU

Tire Pressure Monitor System

- **What is it and what does it do?**
  - Sensors in the tires of vehicles that measure tire pressure
- **How does it work?**
  - ***Show or draw Slide 7: TPMS types and overviews***
  - Direct TPMS systems are connected to the valves of each tire
    - Use pressure sensors, along with ADCs, microcontrollers, batteries, and a radio frequency transmitter to measure tire pressure
  - Indirect TPMS uses software to estimate tire pressure
    - Underinflated tires have smaller radii and higher angular velocities than properly inflated tires

        - These differences could be measured through wheel speed sensors
    - TPMS required by all vehicles manufactured 2008 or later

Radar

- **What does the acronym stand for?**
    - RADAR – Radio Detection and Ranging
    - ***Slide 1: Show and/or draw typical outline of automotive radar***
- **What are the some of the physics behind radar waves?**
    - *** can ask if anybody knows how radars work to begin discussion***
    - As the name suggests, automotive grade radars use radio waves to detect objects in its FOV
    - Radio waves are part of the electro-magnetic spectrum
    - ***Slide 2: Show electromagnetic spectrum, or draw on slide***
    - Automotive radars historically used two bandwidths, 24GHz and 77GHz, but are moving exclusively to 77 GHz
        - For comparison, 24GHz radio waves have a wavelength of around 12.5mm while 77GHz radio waves have a wavelength of around 4mm
        - Typical FM radio waves you use for the radio have wavelength of 2.7 - 3.4 meters (108-88MHz)
        - And visible lights wavelengths vary from 380 – 750 nm (10e-9) at 790 – 405 THz (10e12)
    - The movement to 77GHz allows for larger bandwidths, better resolution, and higher power than the 24GHz band (smaller wavelengths = better resolutions, higher power = longer range)
        - The 77GHz bandwidth actually covers 76-81GHz
- **What are the advantages of radar?**
    - Robust against weather conditions
    - Can operate at night or in bright sunlight
    - Can operate in most cases with moderate snowfall and rainfall
    - Can be mounted behind a bumper
    - Relatively cheap
    - Reliable
- **What type of radars are out there for typical vehicle use?**
    - ***Slide 3: Show or draw typical radar FOV, long range and short range***
    - Automotive radars usually come in 2 types, short range and long range radars
        - Short range radars can have a FOV of around +/- 60-80 degrees out to 100m
        - Long range radars usually have a FOV of around +/- 9-12 degrees out to 250m or more
        - A vertical FOV of around 10-20 degrees is typical for both short range and long range radars
    - Some radars will include both short range and long range in one radar unit

- **What is the cycle time for automotive radars?**
  - Modern automotive radars update anywhere from 14 – 20Hz (50 – 75 ms)
- **So, what type of data do we get from automotive radars?**
  - ***Slide 4: Show or draw typical point cloud for a radar or multiple radars***
  - Raw radar data consists of a point cloud of detections with positional and velocity measurements
  - Each cycle, the radar transmits a signal (Tx) in the form of a radio wave
  - The radio wave spreads out (at the speed of light) and some parts of the wave bounce off objects in the environment and return to mounted radar as the received signal (Rx)
    - Transmitters and receiving antennas are precisely placed to transmit and receive signals only from discrete directions
    - Up to 26 individual antenna nodes
  - The radar is able to differentiate peaks in the Rx signal and determine where those parts originated from in the FOV of the sensor
    - Radars also use the Doppler Shift principle to identify relative speed of each Rx signal
  - The radar can then aggregate and report all the points in the FOV it got an Rx signal from
    - This point can be in polar coordinates (range, angle, elevation, range rate). This is a common 4D measurement for radars.
    - Or be in Cartesian coordinates (x,y,z) with velocity components in the x, y, and z directions
  - The vehicle center or radar mounting point is usually used as the origin for these coordinates
  - Other attributes such as the radar cross section (RCS), and probability of existence can be included in these measurements as well
  - Many cycles can result in up to 100 different detections in the point cloud
    - This can result in around 2,000 points per second
  - Other alternatives to the raw data exist as well
    - Many manufacturers will only have the radar report object tracks and not raw data
    - The radar software will automatically perform filtering and object tracking and report the output as objects with position and velocity estimates
  - The data can be reported over ethernet or the vehicle network depending on the model
- **What are typical radar data statistics?**
  - Weight: ~300 grams
  - Range: 0-250m
    - Range resolution: ~0.5m
    - Range accuracy: ~0.1m
  - Velocity range: -400 – 200 kph
    - Velocity resolution: ~0.5kph

        - Velocity accuracy: ~0.1kph
    - Angle (horizontal) range: +/- 60 degrees
        - Angle resolution: ~1.5-2.0 degrees
        - Angle accuracy: ~0.1 degrees
- **What are some of the specific signal processing and data analysis techniques used in radars?**
    - Frequency modulated continuous wave (FMCW) radars allow the radar to continuously transmit and receive signals while still being able to differentiate which signals come from which cycle
    - Fast Fourier Transforms (FFTs) transfer data from the time-domain to the frequency domain aiding fast signal processing
    - Clustering algorithms can identify data points originating from the same object and eliminate clutter
    - Kalman Filters with extended object tracking capability are implemented in the radar or downstream to take in the raw data, identify certain objects, and then track them over time. This includes automatically adding and deleting objects as they appear within the FOV of the radar
        - Pedestrian, bicycle, vehicle, and truck models can all be used to classify the data in real time
    - *** Slide 5: Show or draw some hardware components for a radar***
    - Specialized multi core CPUs are integrated into the radar's hardware for tasks such as
        - Signal transmission
        - Antennas
        - Signal processing
        - Analog to digital converters
        - Digital to analog converters
        - Power management
- **What is radar used for?**
    - *** ask if anybody knows what ADAS features radars can be used for***
    - ***Slide 6: Show or draw areas of ADAS systems radars can be used for***
    - Radar is one of the most used sensors for ADAS systems
        - Adaptive cruise control
        - Front and rear cross traffic warning
        - Rear collision warning
        - Advanced emergency braking, front collision warning
        - Blind spot detection
        - Automated lane changing
        - Pedestrian automatic emergency braking
        - Automated parking
        - Generating radar-based maps
    - *** Slide 7: Finish up by showing Hyundai radar video on youtube***

Sensor Fusion

- **What is it?**
  - Combining multiple streams of sensor data together to increase knowledge about the state of the vehicle or its surroundings
  - Sensor fusion is a full subject matter area with thousands of publications per year
- **How does it work?**
  - ***Show or draw Slide 1: Sensor Fusion overview and Lidar and camera fusion example***
  - Sensor fusion is actively taking place in all vehicles at different levels and with different subsystems
  - By combining data from different sensors or modalities we can provide the vehicle with a more complete "view" of itself and the environment
  - Sensor fusion aims to combine the strengths of multiple sensors together where a single sensor may have a limitation
    - Radars can detect position and relative velocity but not with as high a resolution and FOV as LiDARs can
    - Lidars can't detect relative velocity like radars can
    - Cameras can detect road signs and other vehicles, but aren't as accurate when estimating position or velocity as radars or lidars are
    - GPS or DGPS can give high accuracy positions of the vehicle itself and a velocity estimate but can lose signal in the city, the mountains or tunnels. IMUs can help keep positional estimates during these times
    - ***Show Slide 2: Short video of radar and camera object detection fusion from NVIDIA***
  - Two main categories of sensor fusion: **centralized** and **decentralized** (also referred to as **low level** and **high-level fusion**)
  - ***Ask if anybody knows the difference between the two approaches***
  - ***Show or draw Slide 3: Low level (centralized) fusion process outline)***
  - Centralized fusion brings together all raw data from sensors to a central computing hub before processing
    - Could be as simple as sending all the wheel speed signal to a central ECU and/or combining it with data from the IMU
      - In fact the ECU is a great example of centralized fusion where inputs from the accelerator pedal, throttle position, steering angle, wheel speed, and transmission all come together to tell the ECU about the state of the car at any given time. The ECU then decides what inputs to give to the engine and other systems
    - Or combining IMU data with GPS data and HD maps for better accuracy
    - Or combining the data returns from multiple radars together to increase the FOV
    - Or include more involved cases, such as combining the individual point clouds from the lidars and radars into one global point cloud
    - Even within sensors like the IMU, the magnetometers can help correct positional heading errors produced by gyroscopic angle drift

- Centralized fusion works best when the sensors have similar data outputs
    - Wheel speed sensors and IMUs can both be used to measure vehicle speed
    - Both lidars and radars contain positional data
    - Multiple instances of the same sensor have the same data output (i.e. stitching different camera frames together)
    - Different data formats must be "aligned" to match up with each other
        - The RGB outputs from cameras need to be comparable with the x,y,z positions of lidar data or radar data
        - May take considerable processing to get them in the same measurement "state space"
        - Asynchronous data coming in from the sensors as well
- Centralized fusion usually needs a powerful central computing system, especially as the amount of incoming data increases (e.g. point clouds)
    - Bandwidth can be restrictive
- The pro of centralized fusion is that all data is shared
    - The camera object detection algorithms can use information from the radar or lidar to cross reference possible object detections and vice versa
    - Regions of interest can be identified in one sensor and used by another to narrow the sensing area to the most important spots
    - Higher levels of certainty when using all the raw data
- De-centralized (high level) fusion is also common
- ***Show or draw Slide 4: High level fusion process outline***
- High level fusion lets subsystems do a lot of the processing and then publishes the results to a central computing system downstream
    - Radars can identify their own objects
    - Lidars can identify their own objects
    - Cameras can identify their own objects
- The central system can then aggregate the object lists coming from each of the sensors into a global object list
- High level fusion is simpler but requires dedicated processing units for each sensor
    - Sensors are larger and more costly due to extra electronics sometimes
    - Introduces more failure modes
    - less maneuverability when designing perception system algorithms than if you were able to use the raw data
    - competing software or hardware architectures from different manufacturers can make implementation harder
- Sensors are not able to talk to each other when performing object detection
    - This can result in less accurate sensing
    - Longer processing times as independent objects from each sensor need to be associated with each other and then classified
- In reality most systems will combine both high level and low level fusion to achieve its sensing needs (referred to as hybrid approaches in some cases)

        - ***Show or draw Slide 5: Hybrid fusion process outline***
        - Some low level sensor data is shared between sensors, such as regions of interest
        - Some sensors perform sensor level tracking and object classification, some may publish data directly to the central processing unit
            - Sometimes feature extraction is done at the sensor level (such as picking out L-shaped corners in camera or radar data) instead of full on object tracking
        - A central processing unit may associate track lists from sensors or add in raw sensor data to the track list
            - May also perform its own high level classifications
        - Outputs from the high level tracking could be shared to the sensors performing their own tracking, as well
        - Pros: all the advantages of low level fusion (global track list feeds back into the sensor level tracks)
        - Cons: more complex systems and algorithms when using a hybrid approach
- **V2X**
    - ***Ask if anybody know what V2V and V2X stands for or how it works***
    - Vehicles can also share information to other vehicles within the vicinity
        - Referred to as V2V
        - Usually involves the sharing of detected objects and vehicle status over radio signal
            - Used to be DSRC (dedicated short range communications) 5.9 GHz band and protocols
            - Is now moving to cellular V2X technology on a more narrow bandwidth around 5.9 GHZ
    - Vehicles can also talk to instrumented intersections or other traffic participants with the technology installed
        - ***Show Slide 6: V2X visualization showing obscured traffic participants at an intersection***
        - V2X
        - Increasing movement to install radars and cameras at urban intersections to track all vehicles and other road users such as cyclists and pedestrians
            - Intersections account for than 50% of crashes each year
        - With the ability to communicate between the vehicle and the instrumented intersection, we can hope to eliminate blind spots as vehicles approach and reduce crashes

Ultrasonic

- **What is it?**
    - ***ask anybody if they know what ultrasonic sensors are or how they work?***
    - Ultrasonic sensors use sound waves to detect objects that are close by, similar to SONAR (also referred to sometimes as proximity sensors or sonar)

- ***Show Slide 1: some examples of ultrasonic sensors***
- They look like dots on vehicles, sometimes hard to notice

- **How does it work?**
  - Small sensors generate sound waves at frequencies greater than 20kHz, typically around 40-70 kHz
    - Obviously higher than the human ear can hear
    - Piezoelectric materials are used to convert electrical signals into the vibrations that produce the ultrasonic wound waves
    - Wavelengths of around 7mm are produced for 48kHz sound
    - Lower frequencies have longer range capabilities due to attenuation
    - Higher frequencies have better resolution but narrower directivity (i.e. FOV will be smaller
    - Ultra sounds can actually emit soundwaves at 100dB or higher, you just can't hear it
  - The sensor then listens for a response, i.e. the sound waves bouncing off nearby objects and returning to the vehicle
    - ***Show Slide 2: Sonar effect and object detection vs geometry***
    - Using time of flight rules and the typical speed of sound through air, distances can be calculated to nearby objects
    - Object detection can depend on the geometry of the object and its orientation
  - Typical automotive ultrasonic sensors can around 80 degree of horizontal FOVs, sometimes up to 120 degrees
  - ***Show Slide 3: Typical ultrasonic sensor configuration****
  - Minimum detection range around 15cm
  - Maximum detection range around 5m, some up to around 10m
  - Update rate, usually 10Hz or higher
  - Vertical FOV is around 50 degrees
    - Ultrasound sensor usually have wider horizontal FOVs than vertical FOVs
  - Most ultrasound sensors only measure 1D, simply range to target
    - Using an array of sensors and more in depth signal processing, you can start to triangulate better positions of objects
    - Advanced systems can also classify objects
    - Many times small objects near the ground are ignored or missed, such as the curb
  - Ultrasonic systems on vehicles almost always include multiple sensors to increase the FOV around the front and rear of the vehicle, if not the whole vehicle
    - ***Show Slide 4: Tesla ultrasound sensors***
    - ***Show Slide 5: Ultrasonic sensor electronics***
    - Use typical electronic subcomponents like ADCs, filters, amplifiers, signal processing
  - Work well in weather and at night in low light conditions
  - Very cheap to make

- **What is it used for?**
    - ***Ask anybody for examples of how ultrasound might be used, if they know***
    - Park assist, parking aid, remote park assist, automated parking
    - AEB for pedestrians and vehicles
    - Blind spot detection
    - Vehicle detection in all directions
    - Tailgate assist
    - Parking guidance systems (tracks open parking spots in garages)
    - ***Show Slide 6: Video of tesla proximity sensors while backing up***