


University of Illinois at Urbana-Champaign
Dept. of Electrical and Computer Engineering

ECE 101: Exploring Digital Information Technologies for Non-Engineers

Autonomous Driving

Autonomous Driving Builds on Earlier Topics

Autonomous driving leverages technologies that we have already discussed:
computer (robot) **vision**,
sensor fusion, and
machine learning.

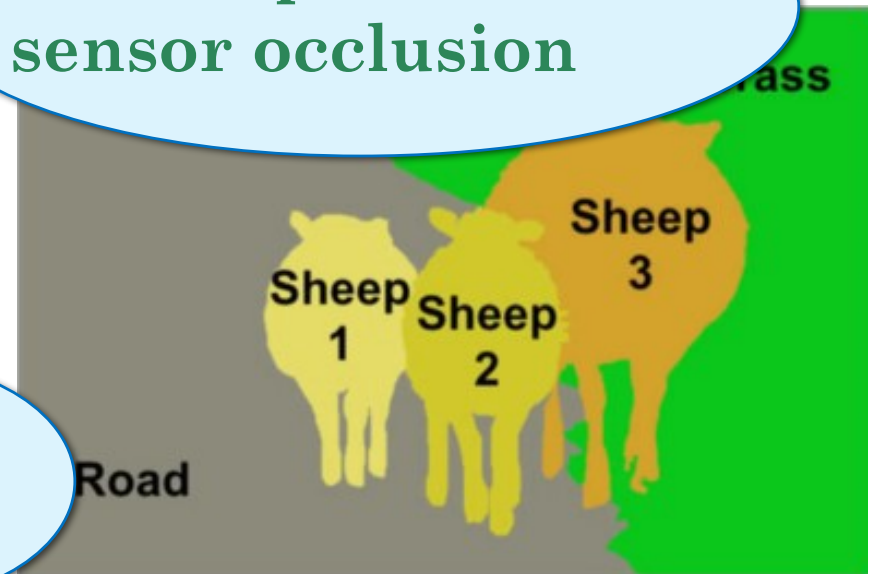


Autonomous Driving Builds on Earlier Topics

Autonomous driving leverages technologies that we have already discussed: computer (robot) **vision**, **sensor fusion**, and **machine learning**.

Can help in case of sensor occlusion

Doing Great

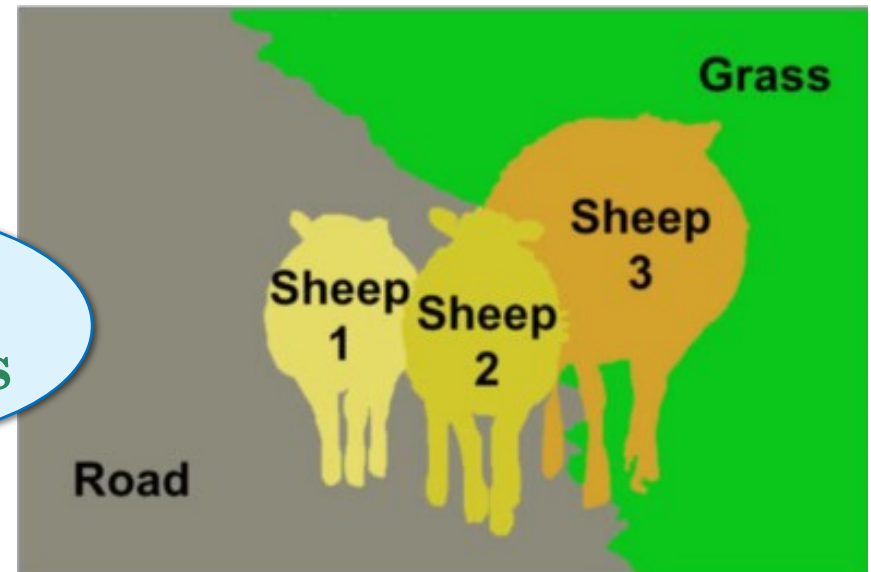


Autonomous Driving Builds on Earlier Topics

Autonomous driving leverages technologies that we have already discussed:
computer (robot) **vision**,
sensor fusion, and
machine learning.



Great
progress



Some Types of Driving Easier than Others

Limited “vocabulary” is also helpful:

driving on freeways is easier than driving in residential areas, which is easier than driving anywhere in arbitrary conditions.



Real Data are Not Easily Acquired for Driving

It's **not that simple**.

Unlike many machine learning applications, **we have relatively little** of the most important types of **data** for training.



Safety Demands Training on the Unusual Events

It's easy and cheap
to pay humans to label digits
or types of clothing.


It's **neither easy nor cheap**
to stage a potential accident
to make sure that autonomous drivers
“learn” to avoid them.

An **autonomous vehicle must be able**
to respond to rare events **safely**.

Companies Leverage Simulation to Generate Data

To address this need, **companies** have **developed** sophisticated **simulations** **that** can **generate sensor data** for a range of physically realistic situations in order **to train** the **ML** models needed **to drive safely**.

Computer games for computers.



ML Models Can be Brittle

That may not be enough, though.


Starting around 2017, studies found that learned **models can be** quite **brittle**.

For example, one model was unable to recognize this decorated stop sign...



Adversarial Approaches are Even More Stark


Adversarial results, in which the models were used to adjust the images, are **even more bizarre**, as illustrated by this ... gibbon.



"panda"

57.7% confidence

+ ϵ



=



"gibbon"

99.3% confidence



Easy to confuse?

Lack of Explainability

A self-driving car suddenly changes lanes on a highway, even though the current lane is clear and safe.

What the system reports: “Lane change executed for optimal route.”

Explainability issue: Engineers cannot determine whether the decision was influenced by:

- A misinterpreted road marking
- A prediction that the current lane would slow down
- A sensor glitch



Can be frustrating for the driver!

Machine Learning is Making Progress on These Problems

Along with lack of explainability,
brittleness to variation and
susceptibility to attack
are **general problems
for machine learning.**

Researchers have been trying
to develop general solutions.

Those will progress.

Sense Compute Communicate Actuate

Let's Focus on the **Actuation** Part of the Cycle


Unlike many of our previous topics, the “**actuate**” part of the cycle is **critical** to **autonomous driving**.

**Quick poll: Who is a better driver?
The human or the self-driving car?**

Three-Point Turns: Rarely Used but Useful

If you know how to drive,
you have probably learned
how to do a 3-point turn.

You probably don't
make many such turns,
but it's a necessary skill
in some situations.



Augmented Reality Adds Computer Graphics to Sight

Of course, real situations often involve obstacles...


**Can you turn around
here without hitting
a car,
a sign,
or a tree,
and without driving
onto the sidewalk?**



Models of Dynamics Express How a Vehicle Moves

Understanding a 3-point turn **requires understanding how your car moves** when you turn the wheels and accelerate.

Models of motion are called **dynamics**, and involve a huge number of factors that we humans usually understand only vaguely.



Humans Have Little Intuition for Dynamics

For example,
humans drive badly in ice, snow, and even light rain because they have **no idea how** these **adverse conditions** **affect** the **friction** between their tires and the road surface.



Models of Dynamics are Quite Sophisticated

More realistic **dynamics** models **incorporate**

- **mass distribution**,
- **acceleration** and **braking**,
- **suspension** and **steering**,
- **aerodynamics**,
- **tires** and **traction** (including issues of slippage both laterally and due to overly rapid braking), and
- even **distortions in** the car's and tires' **shapes** during high-speed turns.

Steps Involved in a Single Driving Operation

Making a 3-point turn involves **several difficult problems**:

- **acquiring a model** of the local environment,
- **selecting the best location** (one that admits a feasible path for the turn and has the fewest safety risks—vehicles traffic, pedestrian traffic, bike traffic, and so forth),
- **path planning based on** the vehicle **dynamics**, and
- **execution of the plan** (including possibly revising or backing out of the plan due to unforeseen complications—for example, someone parks their motorcycle in the space in which your car planned to turn).

Paths Must be Chosen Based on Possible Movement

Path planning is a form of search problem.

In other words, intelligence, as we defined it earlier in the class.

The **constraints** are **imposed by** the vehicle **dynamics**.

For example,

- a vehicle has a turning radius

- which prohibits it from turning too sharply

- (otherwise, we could skip the whole 3-point notion and simply spin the car about its midpoint!).

How Much Distance Required to Stop a Car?

Let's explore a simple example: stopping distance.

A **car** is **driving down a** residential **street** in Illinois. Staying under the speed limit, the car is traveling at **13 m/s (meters per second)**.

How quickly can the car stop?

(You should know if you drive in Illinois!)

https://www.ilsos.gov/publications/pdf_publications/dsd_ds9.pdf

A Simple Formula Allows Calculation of Distance

We'll **use a basic formula for stopping**:

$$\text{distance} = \text{velocity}^2 / (2 \cdot \text{acceleration})$$

Here we assume constant deceleration.

In practice, **deceleration** is **limited by traction**—friction between the tires and the road surface.

Depending on the car and tires,
the limit is **around 7 to 10 m/s²**.

Let's give you a decent car: **8.45 m/s²** deceleration.

Anything Closer than Stopping Distance Demands Choice

$$\text{distance} = \text{velocity}^2 / (2 \cdot \text{acceleration})$$

Plugging in, we obtain...

$$\begin{aligned}\text{distance} &= 13 \cdot 13 / (2 \cdot 8.45) \\ &= 169 / 16.9 \\ &= 10 \text{ meters (33 feet, 11 yards)}\end{aligned}$$

That's assuming an instantaneous reaction.

If anything gets into the next 10m of the car's path, either **the car has to swerve or hit the object.**

Higher Speeds and Lower Traction Increase Distance

What if instead you're driving on a country road?

The speed limit there is about **24 m/s (~50 mph)**.

There's also some gravel on the road
to protect against ice in winter,
so maximum deceleration is a bit lower: **7.2 m/s**.

Now how much space does the car need?

The Result? Four Times as Long...

$$\begin{aligned}\text{distance} &= 24 \cdot 24 / (2 \cdot 7.2) \\ &= 576 / 14.4 \\ &= 40 \text{ meters (131 feet, 44 yards ...} \\ &\quad \text{almost half of a football field!)}\end{aligned}$$

That's a long way! Hope no deer are nearby.

Multipliers for Various Surfaces May Surprise You

How do adverse conditions affect stopping distance?

The table gives examples relative to dry asphalt.

(Note that our model is simple. Friction and therefore deceleration goes down with higher velocity, but drag from air goes up.)

Surface	Relative Deceleration	Relative Stopping Distance
Dry asphalt	1.00	1.00
Wet asphalt / gravel / sand	0.667	1.50
Dry earth	0.755	1.32
Wet earth	0.611	1.64
Packed snow	0.222	4.50
Ice	0.111	9.00

Right off “Illinois Rules of the Road”

Doubling your speed **quadruples** your vehicle’s stopping distance.

Consider the following when deciding your speed:

- How quickly you can react physically and mentally.
- Type and condition of the roadway.
- Size of the tires — large, wide tires with good tread will stop a vehicle faster than small, narrow tires with little tread.
- Condition of the brakes — newer brakes stop a vehicle more quickly than older, worn brakes.
- Direction of the wind and how fast it is blowing — a strong tail wind can make it very difficult to stop.
- Type of vehicle — vehicle design, weight distribution, suspension and shock absorbers all play a role in how quickly a vehicle can stop.

“Visual” Data Fused with Other Types of Data

“Visual” data are likely to be **integrated with other types** of information.

Let’s say that my car is following another vehicle, for example.

How do we know when the vehicle in front slows down?

1. **Brake lights**
2. **WiFi notification** from vehicle (in future)
3. “Visual” **speed/depth change** detection



Most Sensors Must be Treated as Unreliable

1. Brake lights
2. WiFi notification from vehicle (in future)
3. “Visual” speed/depth change detection

The **first two may fail:**

Sensor occlusion

- slowing may occur without braking, and
- brake lights can go out or be obscured.
- WiFi notifications may not arrive, or may not be sent from older cars.



Trust What You See?

1. Brake lights
2. Notification from vehicle (in future)
3. “Visual” speed/depth change detection

Only change detection is reliable

- (and even there, we need to have multiple sensors—ours can also fail),
- but it is **also delayed**—takes time to notice.



Safety-Critical Systems Require Safety Envelopes

In practice, systems require a “**safety envelope**”:

- make sure that the car
- has **enough time to react safely**
- **to changes** in the environment.



Safe Decision Making Relies on Layers of Models and Data

Information on slowdown

- must be fed into **models for motion planning**,
- which in turn rely on accurate **models of vehicle dynamics**,
- which rely on sensing the **vehicle's load**.

Average car weight is 4,000 pounds,
with a cargo load of about 1,000 pounds.

That's a 25% change in mass,
assuming humans obey rules.



How Safe is Safe? Who Defines Autonomous Behavior?

How safely should an autonomous car drive?

All US schools teach defensive driving:

assume that other drivers are going to make mistakes.

Humans Don't Assume Crazy Drivers

What kind of mistakes?

If I'm driving

- in the right lane of a **four-lane road** (two in each direction), and
- an oncoming vehicle is in the far lane (the right lane on their side),
- **should I assume that they might swerve in front of me at any time?**

Probably not.

Humans Actually Usually Assume Good Drivers

What if it's a two-lane road?

Head-on crashes

- due to drunk driving,
- inattention (texting), and so forth
- are much more likely to lead to serious injury.

Maybe just don't drive on two-lane roads?

The Right Answer is Hard to Define

Realistically, like humans,

- autonomous driving should fall **somewhere between**
- **timid** (slow down, there's a car coming!) **and**
- **oblivious** (so what? it's MY turn to use that part of the road to pass!).

It's easy to say that **both extremes are bad**.

Exactly **how** the **car should behave** is **not so easy** to specify.

Need Legal Standards and Behavioral Expectations

Accidents will happen,
even if all vehicles are autonomous
(perhaps rarely if cars without
new tires and an oil change),
and **people will die.**

We **need legal standards for safety**
and **expectations for behavior.**

Lack of explainability in AI won't help.


Careful statistical comparison with human drivers may.

Safety Devices for Vehicles

Example: Safety device added to trains in mid-19th century!

The red grill in front of this train engine was invented by Charles Babbage in 1838, about 16 years after he invented the programmable computer.

Its common name reveals its purpose:
cow catcher



Cow Catcher Ensures that Train is Safe, But Not Cow

A cow on train tracks might be pushed down under the engine's wheels, derailing the train.

With a cow catcher, the cow is flung to the side (and invariably killed) without damaging the engine!



To Swerve or Not to Swerve?

You're driving on a narrow mountain road.

Suddenly, someone runs down the hill onto the road.

You don't have time to brake!

Do you swerve off the road and over the cliff, **or run Pat down?**

Does it depend how many family and friends are with you in the car?



Whose Life Comes First?

**What's more important,
pedestrian's life or passenger safety?**

If a car has to decide between hitting a pedestrian and endangering the vehicle (and thus the passengers), what should it do?

Many questions that need to be answered.

Issues that Persist

Technical Difficulties

- **Perception challenges:** Handling poor weather, low light, and sensor occlusion.
- **Edge cases:** Reacting to rare, unpredictable scenarios like unusual road hazards.
- **Predicting human behavior:** Anticipating erratic actions of drivers and pedestrians.
- **Data and computation:** Requires massive datasets and real-time processing power.
- **Infrastructure variability:** Roads, signs, and traffic systems differ across regions.

Issues that Persist

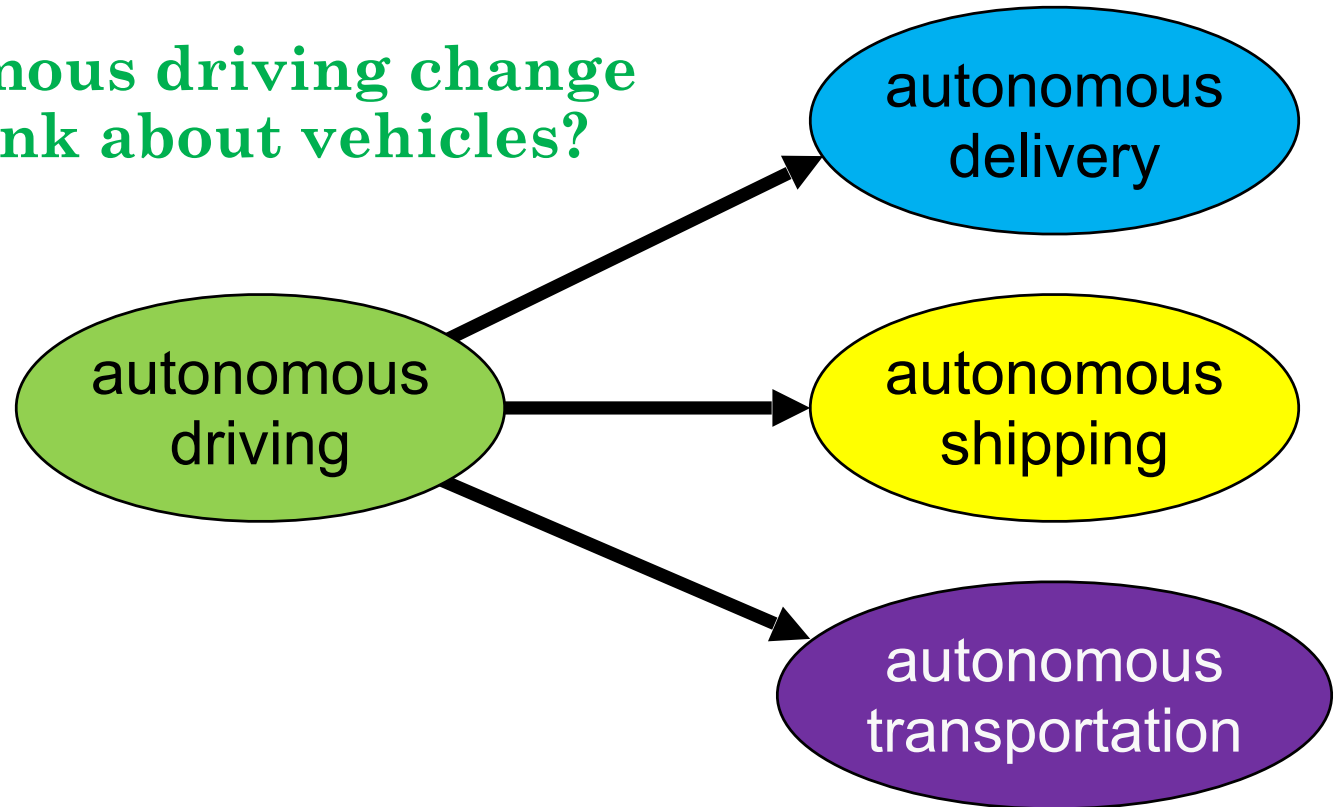
Ethical Difficulties

- **Crash dilemmas:** Deciding whom to prioritize in unavoidable accidents.
- **Transparency:** Explaining why the car made a specific decision (lack of explainability).
- **Liability:** Determining who is responsible in case of accidents —manufacturer, software developer, or owner.
- **Privacy concerns:** Managing sensitive data collected by sensors and cameras.

Autonomous Driving Enables Other Uses

Will autonomous driving change how we think about vehicles?

Possibly ...



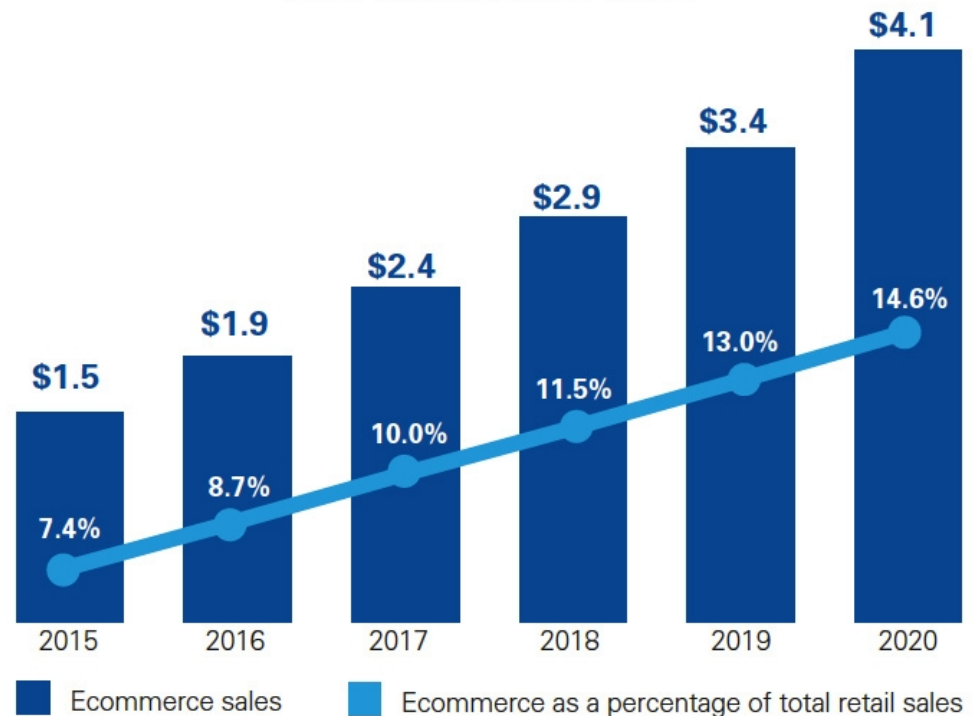
Autonomous Driving Enables Autonomous Delivery

autonomous
delivery

Online sales account for nearly 1/6th of all sales.

Autonomous driving enables autonomous delivery of online purchases.

Online retail sales as a percentage of total retail sales (US\$ trillions) 2015-2020



Autonomous Shipping: Optimization of Supply Chains

Autonomous driving enables autonomous shipping (trucks and trains).

Distribution of goods

- based on average consumption
- adjusted for variations
- in online shopping demand.

Large chains can integrate

- **from inventory control**
- **through distribution** all the way
- **to ordering** from suppliers.



autonomous
shipping

Autonomous Driving May Enhance Public Transportation

Transportation rental companies

- such as Uber, Lyft, and so forth
- have become **popular internationally**
- for everything from vans to scooters.



autonomous
transportation

Autonomous driving enables

- these **services to be automated and**
- to be **optimized** for efficiency,
- perhaps overcoming cultural barriers
to public transportation.

Terminology You Should Know from These Slides

- autonomous driving
- sensor occlusion
- vehicle dynamics model
- path planning
- stopping distance
- autonomous delivery
- autonomous shipping
- autonomous transportation

Concepts You Should Know from These Slides

- why simulations are necessary for training autonomous vehicles
- why general ML problems such as brittleness, vulnerability, and lack of explainability are more important when safety is an issue
- aspects covered by vehicle dynamics models
- steps in a driving operation: acquiring environment model, selecting a location, path planning, plan execution
- how to calculate stopping distance and why it matters for driving
- why defining an acceptable safety level is difficult
- how autonomous driving might change how we think about delivery, shipping, and transportation