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

ECE 101: Exploring Digital Information Technologies for Non-Engineers Fall 2024

Exam 3 Review

Theme: Sense-Compute-Communicate-Actuate Loop

- **Drone delivery systems**
- **Self-driving cars (autonomous vehicles)**
- **Farm robots**
- **Digital assistants:** Alexa, Siri, etc.
- **Cleaning Robots**: Roomba
- **Smart Treadmill**
- **Smart water bottle**
- …

Example: Drone Delivery System

- **Sense:**
  - Sensors for temperature, pressure, wind speed—environment conditions,
  - GPS—location
  - Camera, IMU
- **Compute**:
  - Process input—sensed data coming in as input,
  - Provide output—decision on what to do
- **Communicate**:
  - Send updates to and get instructions from control station
  - Communicate with other drones (if part of a fleet)
- **Actuate:**
  - Control rotors to land safely or hover—speed up/down, change directions etc.
  - Control mechanism to deposit package safely.

Example: Autonomous Vehicles

- **Sense:**
  - Camera
  - LiDAR
  - IMU, GPS, Microphone
- **Compute**:
  - Process input—
  - Provide output—
- **Communicate**:
  - With the human in the car
  - With a centralized control system
- **Actuate:**
  - Control engine, speed up/down, turn wheels, apply brakes, turn signal
  -

Types of Sensors

Types of sensors

1. **Cameras**: Regular cameras, IR, thermal, radar, Lidar
2. **Microphones**: audible, ultrasound
3. **IMU**: **I**nertial **M**easurement **U**nit (accelerometer, gyroscope, magnetometer)
4. **Wireless**: GPS, Wifi (WiGig 60GHz, THz), UWB
5. **Assorted**: pressure, humidity, proximity, temperature, chemical traces

Basic working principle

- **Pin-hole camera**
- **Microphone**
- **LIDAR**
- **GPS**

Computer Vision: Basic Image Processing Techniques

- **Edge detection**
- **Image segmentation**
- **Feature extraction/registration**
- **Depth perception**

Computer Vision: Things to remember

- Active vs. passive sensors … why do we need both?
- Sensor fusion
- Occlusion
- Safety envelope

Speech and NLP: Issues

- Context
- Ambiguity
- Sarcasm, irony, euphemism, metaphors
- Accents, dialects, pidgin languages etc.

Computing steps to understand speech

1. **get rid of noise**: other voices (unauthorized), music, television, video games, pets, and so forth.
2. perform "**voice recognition**": translate an audio signal into a sequence of words.
3. understand what the human is trying to communicate: **process their natural language** (English, for example).

Computational tools for understanding speech

1. Context— search for keywords
2. Share information between hierarchy of interacting, probabilistic models
3. MLE - Maximum Likelihood Estimate: Given an observation,
   - **choose** the **explanation** that is
     - **most likely to produce** the **observation.**
4. Bayes' Theorem:

   Probability (A AND B) = Probability (A) x Probability (B | A)

   The chance of **A and B both happening** is equal to **the product of the chance of A happening and the** chance of **B happening if A has happened**.

A game of probability

Pat will roll either

- one (six-sided) die or
- two dice and add up the numbers.

Then Pat tells us the amount rolled. Can we guess whether Pat rolled one or two dice?

Probability Pat rolls 1 die = 1/2 (one of two options they have)
Probability Pat rolls 2 dice = 1/2 (one of two options they have)

What is the chance that Pat rolled a 4 with 1 die? 1/6
What is the chance that Pat rolled a 4 (total) with 2 dice? 3/36 or 1/12
(Could have been 1+3, 2+2, or 3+1)

probability **(Pat rolled one die)** X probability **(got a 4 | Pat rolled one die)** = 1/2 x 1/6 = 1/12
probability **(Pat rolled one die)** X probability **(got a 4 | Pat rolled one die) =** 1/2 x 1/12 = 1/24

Initial Probabilities are Important to Correct Choices

What **if Pat tells us** more about how they choose to roll 1 or 2 dice?

probability (Pat rolls one die) = ¼ and

probability (Pat rolls two dice) = ¾

In that case, **our guess changes**, as

¼ · probability (got a 4 | Pat rolled one die)

= 1/4 x 1/6 = 1/24

<

¾ · probability (got a 4 | Pat rolled two dice)

= 3/4 x 1/12 = 1/16

How is MLE used?

Voice recognition answers the question, **"Given an audio input, what sequence of words was spoken?"**
A solution is generated by finding the sequence of words that is most likely to have generated the audio input.

Natural language processing answers the question, "**Given a sequence of words, what did the speaker want to communicate?**"
A solution is generated by finding the meaning that is most likely to have generated the sequence of words.

BERT: Bidirectional Encoder Representations from Transformers

Guessing Words Easier with Words on Both Sides

"I took my … [more words]."
What's the next word?
In 1953, journalists* realized that **the words AFTER the missing would help** in guessing.

"I took my … for a walk."
What's the next word?
Dog, perhaps?
Could be lots of words, but dog may be a good MLE choice.
"Cat," "snake," "Ferrari," maybe not so good.

BERT: Bidirectional Encoder Representations from Transformers

BERT (**Bidirectional Encoder Representations from Transformers**)
○ revolutionized how machines understand language
○ **read text in both directions at once**, which helps it understand the context better (before BERT, many models read text in one direction i.e. left to right)
○ first trained on a large amount of text (**pre-training**) to learn general language patterns
○ **fine-tuned** to the specific task, like answering questions or sentiment analysis (versatility)

Large Language Models (LLMs)

GPT (Generative Pre-trained Transformer), developed by OpenAI.

1. Generative Model: While BERT is great at understanding and processing text, GPT is designed to generate text. This means it can write essays, create dialogue, and even compose poetry.

2. Transformer Architecture: Both BERT and GPT use a type of neural network called a Transformer, a neural network that learns context and thus meaning by tracking relationships in sequential data like the words in this sentence. GPT uses this architecture to predict the next word in a sentence, making it very good at generating coherent and contextually relevant text.

3. Pre-training on a Massive Scale: GPT models are pre-trained on a vast amount of text from the internet. This extensive training helps them generate more natural and diverse text.

Applications of NLP (Generic or Specific)

- **Text Classification:** Categorizing text into predefined categories.
  - Did the reviewer like a movie?
  - What sentiment did they express?
- **Interpretation:**
  - Is anything in a patient's electronic medical records relevant to the patient's current symptoms?
- **Named Entity Recognition (NER)**
  Identifying and classifying entities like names, dates, and locations in text.
- **Question Answering:** Finding answers to questions within a text
- **Text Generation**: Creating new text based on a given prompt.
- **Dialogue Systems**: Powering chatbots and virtual assistants.
- **Creative Writing**: Generating stories, poems, and other creative content.

AR, VR, MR and ER

- **AR - Augmented reality**, designed to **add digital elements over real-world views** with limited interaction.
- **VR - Virtual reality** immersive experiences helping to **isolate users from the real world**, usually via a headset device and headphones designed for such activities.
- **MR - Mixed reality** combining AR and VR elements so that **digital objects can interact with the real world**, means businesses can design elements anchored within a real environment.
- **ER -Extended reality (XR)**— covering all types of technologies that **enhance our senses**, including the three types previously mentioned.

AR/VR Applications

- **Ones discussed in class**
- **New ones**

AR/VR Helpers

- **Motion Capture**
  - With markers and marker less
  - Applications
- **Haptics**
  - Whats is it?
  - Applications

Autonomous Driving: Major Technologies

- computer (robot) vision,
- sensor fusion, and
- machine learning

Autonomous Driving: Issues

- Unlike many machine learning applications, we have **relatively little** of the most important types of **data** for training.
  - 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**
- **Failed** image recognition
  - ML Models Can be Brittle
  - ML models can fall prey to adversarial attacks leading to

Autonomous Driving: Data Scarcity Solution

- **S**ophisticated **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)
- Collect more data (millions of Teslas on the road)

Autonomous Driving: Actuation

- **Examples** of how actuation needs of a self driving car lend themselves to great application of computation
  - 3 point turn
  - Driving on different surfaces
  - Calculation of stopping distance

Autonomous Driving: Safety Concerns

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.

Autonomous Driving: Applications, Advantages and Disadvantages

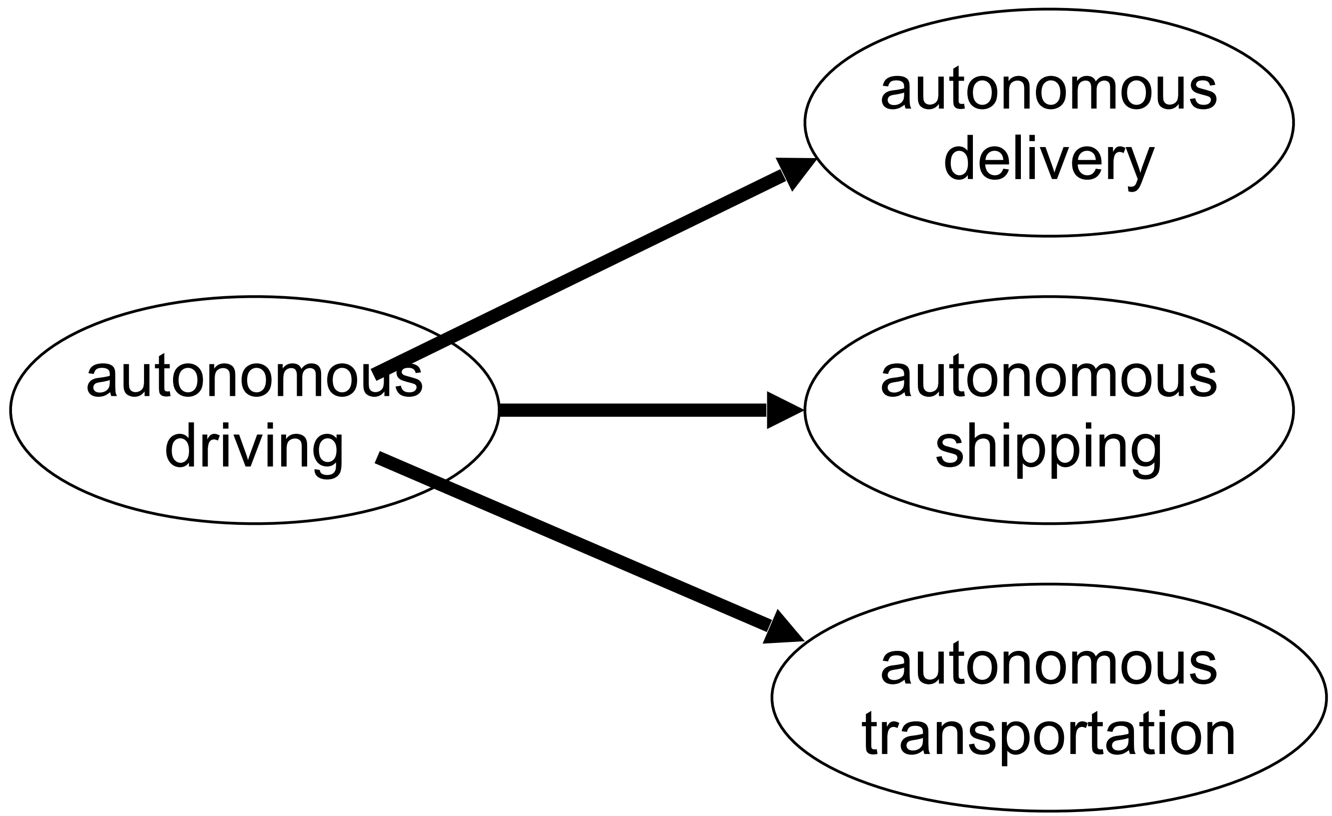

File Systems

- **Cloud service:** Desired properties: Availability, Reliability, Consistency
- Undoing operations: Undo log, snapshots, version control
- Why does order of operations matter?
- Why might inconsistency occur in a shared file? How is it resolved? (Server version that serializes operations)
- Why would we tolerate inconsistencies? For Speed

File Systems

- **Eventual consistency**
- **Push and pull models**

Last Class Participation Assignment

- **Complete ICES survey and confirm on Canvas.**
  - Abrita
  - Sattwik

  - **Your feedback matters! I would like to make the course more useful to future students.**