CS598JHR Embodied NLP

http://courses.grainger.illinois.edu/cs598jhr

Welcome to CS598JHR Embodied NLP!

Julia Hockenmaier

juliahmr@illinois.edu

Slides Version: 1/24/24

ILLINOIS

How will we teach this class?

Who we are

Instructor: **Julia Hockenmaier (Julia)**
Professor, Computer Science
Siebel 3324
juliahmr@illinois.edu

TA: **Risham Sidhu**
PhD student, Computer Science
rsidhu3@illinois.edu

What is the purpose of this class?

Graduate seminar class on "Embodied" NLP

Focus on **paper reading/presentations** and **research**

We assume you have taken a general NLP class

(e.g. CS447 or CS546)

Graduate seminar class on **"Embodied" NLP**

"Embodied": Agents that operate in a real or simulated environment

NLP: Natural Language Processing (i.e. Understanding, Generation and Dialogue)

"Embodied" NLP: Language understanding, generation and dialogue for/with embodied agents

This class consists of…

… lectures

Wednesdays/Fridays,11:00am—12:15pm, 2200 Sidney Lu

Many of these will be paper presentations by students

… office hours

TA office hours are intended for hands-on help with projects

My office hours are mainly intended for paper presentations

… research projects

These can be done in groups of up to four students

… a Canvas page

For grades and to submit reports and paper reviews

… a Campuswire page

For discussions and to find teammates for projects

… a website https://courses.grainger.illinois.edu/cs598jhr

For slides, syllabus etc.

Assessment

Your grade will consist of

… 35%: your presentation of a research paper in class

… 50%: your research project

… 10%: your written reviews of research papers
(graded mostly for completion)

… 5%: your participation in class

Paper presentations

Everybody needs to prepare a 15-minute oral presentation and a two-page writeup about one research paper to be shared with the class.

NB: This paper shouldn't come from your own research group, nor can it be a paper you presented in your qualifying exam.

- We will send out a sign-up sheet with dates and papers for each class.
- You will have to come to my office hours *the Monday or Wednesday of the week when you're presenting* with your slides to show them to me, otherwise you will only get half credit for your presentation.
- You have one week after your presentation to send in your writeup (so that you can reflect any in-class discussion)

Short Paper Reviews

For 10 of the papers that we discuss in class, you will have to submit a short review before that class

- —Due to the size of the class, we can largely grade you for completion (although we will spot-check your answers)
- — You will have to submit the reviews through Canvas.
- — In the past, we've used a LaTeX template for this, but we may switch to tests inside Canvas

We encourage you to get into the habit of taking notes about the papers you read. Hopefully this will get you started!

Paper presentations: procedure

Presenter:

- **Meet with me** at least two days before your presentation

We want to make sure you understand the paper

- **Slides** are recommended, but: please **make your own**, even when the authors make theirs available

You don't actually learn much by regurgitating somebody else's slides.

- Send me a **PDF of your slides before class**
- Bring your laptop (or let me know in advance if you need to use mine)

Everybody else:

- Before class: submit a **one-page summary of the paper**

I won't grade what you write, but I want you to engage with the material

- During/after class: **critique** the presentation

This is merely for everybody's benefit, and not part of the grade.

Research projects

You will have to complete a sizable research project.

Due to the size of the class, you will have to work in groups (we're aiming for 3–4 students/team).

There will be several milestones:

- Initial proposal
- Intermediate report and presentation
- Final report and presentation

Research projects

The aim is for each team to produce something that could be submitted to a conference:

— You should aim to make an actual contribution to research

— Your presentation should be sufficiently polished

If you build on existing research, talk to me, and loop your advisor in as well if necessary.

If you're doing related projects in other classes, let me and the other professor know.

DRES accommodations

If you need any disability related accommodations, talk to DRES (http://disability.illinois.edu, disability@illinois.edu, phone 333-4603)

- If you are concerned you have a disability-related condition that is impacting your academic progress, there are academic screening appointments available on campus that can help diagnosis a previously undiagnosed disability by visiting the DRES website and selecting “Sign-Up for an Academic Screening” at the bottom of the page.”

Come and talk to me as well, especially once you have a letter of accommodation from DRES.

- Do this early enough so that we can take your requirements into account

What will we cover in this class?

How does embodiment affect NLP?

People learn and use language
in the physical world

People give and follow instructions
in the physical world

People use language to collaborate
in the physical world

Most language in the physical world is spoken,
rather than written

How do we teach NLP in 2024?

Research is happening so fast that it is difficult for anybody to keep up

Your suggestions of what papers to read are most welcome

LLMs are revolutionizing NLP/AI.
To do NLP now, you need to understand LLMs.
But NLP is still much more than LLMs.
To do NLP *well* now and in the future,
just knowing LLMs is not sufficient

How does embodiment affect NLP?

Communication in/about the physical world requires the ability to map ("ground") linguistic references to the physical world

The environment we communicate in/about may be changing during a conversation

Different people/agents may perceive the same environment/situation differently. Successful communication requires communication partners to adapt to each other's perceptions.

How does embodiment affect NLP?

Physical and simulated environments and tasks vary a lot

Sensors and actuators of embodied agents vary a lot

In practical terms, we have much less environment-specific training data for embodied tasks