

---

# Natural Language Processing

## Conclusion and Q&A

**Sofia Serrano**  
**[sofias6@cs.washington.edu](mailto:sofias6@cs.washington.edu)**

Credit to Noah Smith for slides

# Announcement(s)

- A3 is due tonight at 11:59pm!
  - Don't forget to also submit your \*.preds files and your writeup pdf (and tag your A3)!
    - Not sure whether those have been uploaded correctly? Check <https://gitlab.cs.washington.edu/cse447-wi23/a3/cse447-wi23-a3-yournetid/> in a browser.
  - The number of late days you have available for A3 is  $\min(3, 5 - \text{total\_late\_days\_used})$ 
    - Not sure what this is for you? Ask us!
  - We'll be taking A3 regrade requests from whenever we get A3 grades out next week to the end of next week's Friday (so, probably for 2ish days)
- Please please please do us a favor and **fill out the end-of-course eval sometime before this Sunday (3/12) at 11:59pm!**
  - The end-of-course eval is our best tool for figuring out what to keep the same, and what to change, for future iterations of this course
  - Future 447 students benefit from your feedback!!

**A high-level view of this  
quarter**

---

# What the syllabus has ended up as

- **Introduction**
  - Overview of NLP as a field
- **Modeling (ML fundamentals)**
  - Text classification: linear models (perceptron, logistic regression), non-linear models (FF NNs, CNNs)
  - Language modeling: n-gram LMs, neural LMs, RNNs
  - Representation learning: word vectors, contextualized word embeddings, Transformers
- **Linguistic structure and analysis (Algorithms, linguistic fundamentals)**
  - Words, morphological analysis (this last one not so much– if you're curious about this topic, see [Noah's 517 lecture/slides](#) on this topic)
  - Sequences: part of speech tagging (POS), named entity recognition (NER)
  - Syntactic parsing (phrase structure, dependencies)
- **Applications (Practical end-user solutions, research)**
  - Toxic language detection, commonsense and socially-informed reasoning
  - Multimodal NLP
  - Prompting and in-context learning with large language models
  - Multilingual NLP

# Recurring themes

- The role of machine learning, and the evolution of techniques (relative frequencies, log-linear models like logistic regression, neural networks)
- The role of data (e.g., annotation, parallel data)
- Challenges specific to natural language (ambiguity, variation in language)
- The importance and challenge of evaluation
- Useful building blocks, from high-level abstractions (e.g., noisy channel) to low-level tools (e.g., transformers, dynamic programming)

# A grab bag of some of the many many topics we didn't have time to cover

- Discourse and pragmatics
- Topic modeling (a tool originating in NLP that's now commonly used in some social science research areas)
- Other popular NLP tasks: Natural Language Inference, Question Answering, Summarization
- (More on) graphical models, [structured prediction](#)
- [Model interpretability](#)
- “Deeper” coverage of neural networks (e.g., theory, more advanced learning algorithms)
- Speech processing

# So why did we structure the syllabus this way?

In short, to expand your “zoo” of NLP methods that you’ve seen.

With this set of content, a wide range of different problems in NLP will either (a) look very similar to something we covered, or (b) tend to have approaches that look at least similar to approaches we’ve walked through.

# Things we'd like to be true about our NLP methods

1. Sensitive to a wide range of the phenomena and constraints in human language
2. General across different languages, genres, styles, and modalities
3. Computationally efficient at construction time and runtime
4. Have strong formal guarantees (e.g., convergence, statistical efficiency, consistency, etc.)
5. Highly accurate, as judged against expert annotations and/or task-specific performance
6. Explainable to human users

We are often missing some of these!  
But they're still good to have as goals.



# Approaching NLP problems: A rough template

---

# What's my *problem*?

- What do I want to do?
  - e.g., Create a system that makes automated decisions about text? Conduct an analysis of a large dataset?
  - Which problems have I seen in the past that are similar? Are there any fundamental differences between this problem and ones I've seen?
- What **data** relevant to my problem do I have available?
  - Ethical considerations for this particular data?
  - How much data do I have?
  - What kinds of labels (if any) are included with this data?
  - Are there other, existing resources that aren't quite data for my problem but could still somehow be useful (e.g., other data in the same language, or files of pretrained parameters for models/word embeddings)?
  - Do I have the resources to create data for this problem?

# What kind(s) of modeling approach(es) will I try?

- SUPER related to the questions about related tasks and amount of data/kinds of labels in data
- Can I do fully-supervised training? Or does my data force me to do something semisupervised or unsupervised?
- How many model parameters could I reasonably learn with this amount of data?
- Am I planning on learning parameters via some variation on gradient descent? If so, what do I want my loss function to minimize?

# How will I evaluate my approach?

SUPER important.

In this class we talked about canonical evaluations for the tasks we covered, *but those still have problems!*

Some questions to consider:

- What kinds of blind spots is my planned evaluation likely to have?
  - Often related to “what do(es) my metric(s) of choice miss?” or “what are the limitations of my test data?”
- How well does my test data represent the cases I have in mind that I’m hoping my model will generalize to?
- Does some mixture of human and automatic evaluation make sense/fall within my available resources?
- Do I want to evaluate my model on the task I trained it to perform, or evaluate it by looking at how it helps another downstream model perform a different task?

# Resources to keep learning



# At UW: Classes

- CSE 472 ([sample past syllabus](#)) and CSE 517 ([sample current syllabus](#)) overlap with this course a *lot*, but worth checking their syllabi to look over content we didn't cover (and if that's the content that interests you most, then maybe worth taking)!
- CSE 481N (the NLP capstone, [sample past syllabus here](#)– Noah's running this next quarter!) is a great way to dive into a slightly longer-term NLP project on a topic you're interested in
- If you've taken AI: [CSE 582 Ethics in AI](#) is a new course being taught by Yulia next term!
- In general, within CSE, keep an eye out for courses taught by Luke Zettlemoyer, Yejin Choi, Noah Smith, Hannaneh (Hanna) Hajishirzi, Yulia Tsvetkov
  - Their 599s typically focus on reading and discussing papers

# At UW: Not-classes

- #uwnlp channel on [UW CSE slack](#) (have a CSE email? Then you can join without help– otherwise, feel free to tell me to add you!), a collection of a *lot* of NLP people within CSE
- Mailing list for announcements of NLP talks at UW:  
<https://mailman.cs.washington.edu/mailman/listinfo/uw-nlp>
- Research (if you'll be around UW for at least a couple more quarters)
  - Different labs within CSE have different protocols for how to get involved. Examples:
  - [Noah's ARK](#) (“If you would like to apply to work in Noah's ARK”)
  - [H2Lab](#) (Hanna's group: “For Undergraduate Students: Applying to work at H2Lab”)
- (Date is tentative) Next quarter, 2nd-week Wednesday (April 5), 5:30-8pm:  
UW-wide panel on ChatGPT with a bunch of CSE NLP faculty, plus faculty from other departments who are thinking about implications of ChatGPT for education

# Beyond UW

- [NLP Highlights](#) podcast
- Want to learn some new NLP techniques?
  - Is what you're looking for, or something close, mentioned in [the J&M textbook](#)? If so, start there
  - If it's a popular topic, often there are nice explanatory blog posts online!
  - Where/how to find NLP research papers:
    - Just browsing for interesting stuff?
      1. Go to the ACL anthology: <https://aclanthology.org/>
      2. In the top table on the page ("ACL Events"), click on a recent year number by a venue (AAACL, EACL, NAACL, TACL, EMNLP, Findings are the broadest)
      3. In the (probably massive) list of paper titles on the resulting page, browse or ctrl-F for a specific keyword in a title
    - Looking for a specific topic and want better search than the built-in one on ACL anthology?
      - <https://www.semanticscholar.org/>
- Other non-ACL venues that might be of interest: [ICLR](#), [ICML](#), [NeurIPS](#), [Text as Data](#), [AAAI](#), [ICDAR](#), IEEE journals/conferences for speech processing...



Thanks for a great quarter everyone! :)  
Questions?

(I'll just casually drop the [end-of-course eval link](#) again...)