Ling/CSE 472:
Introduction to Computational Linguistics

3/28/31

Introduction, overview

Who’s here?

- A good class to work together --- everyone brings different skills
- I’m going to bring a lot to this class because...
- This is going to stretch me because...

Current events - some polls about ChatGPT

- All should allow anonymous participation
- I’m interested in your honest opinion!

ChatGPT is AI

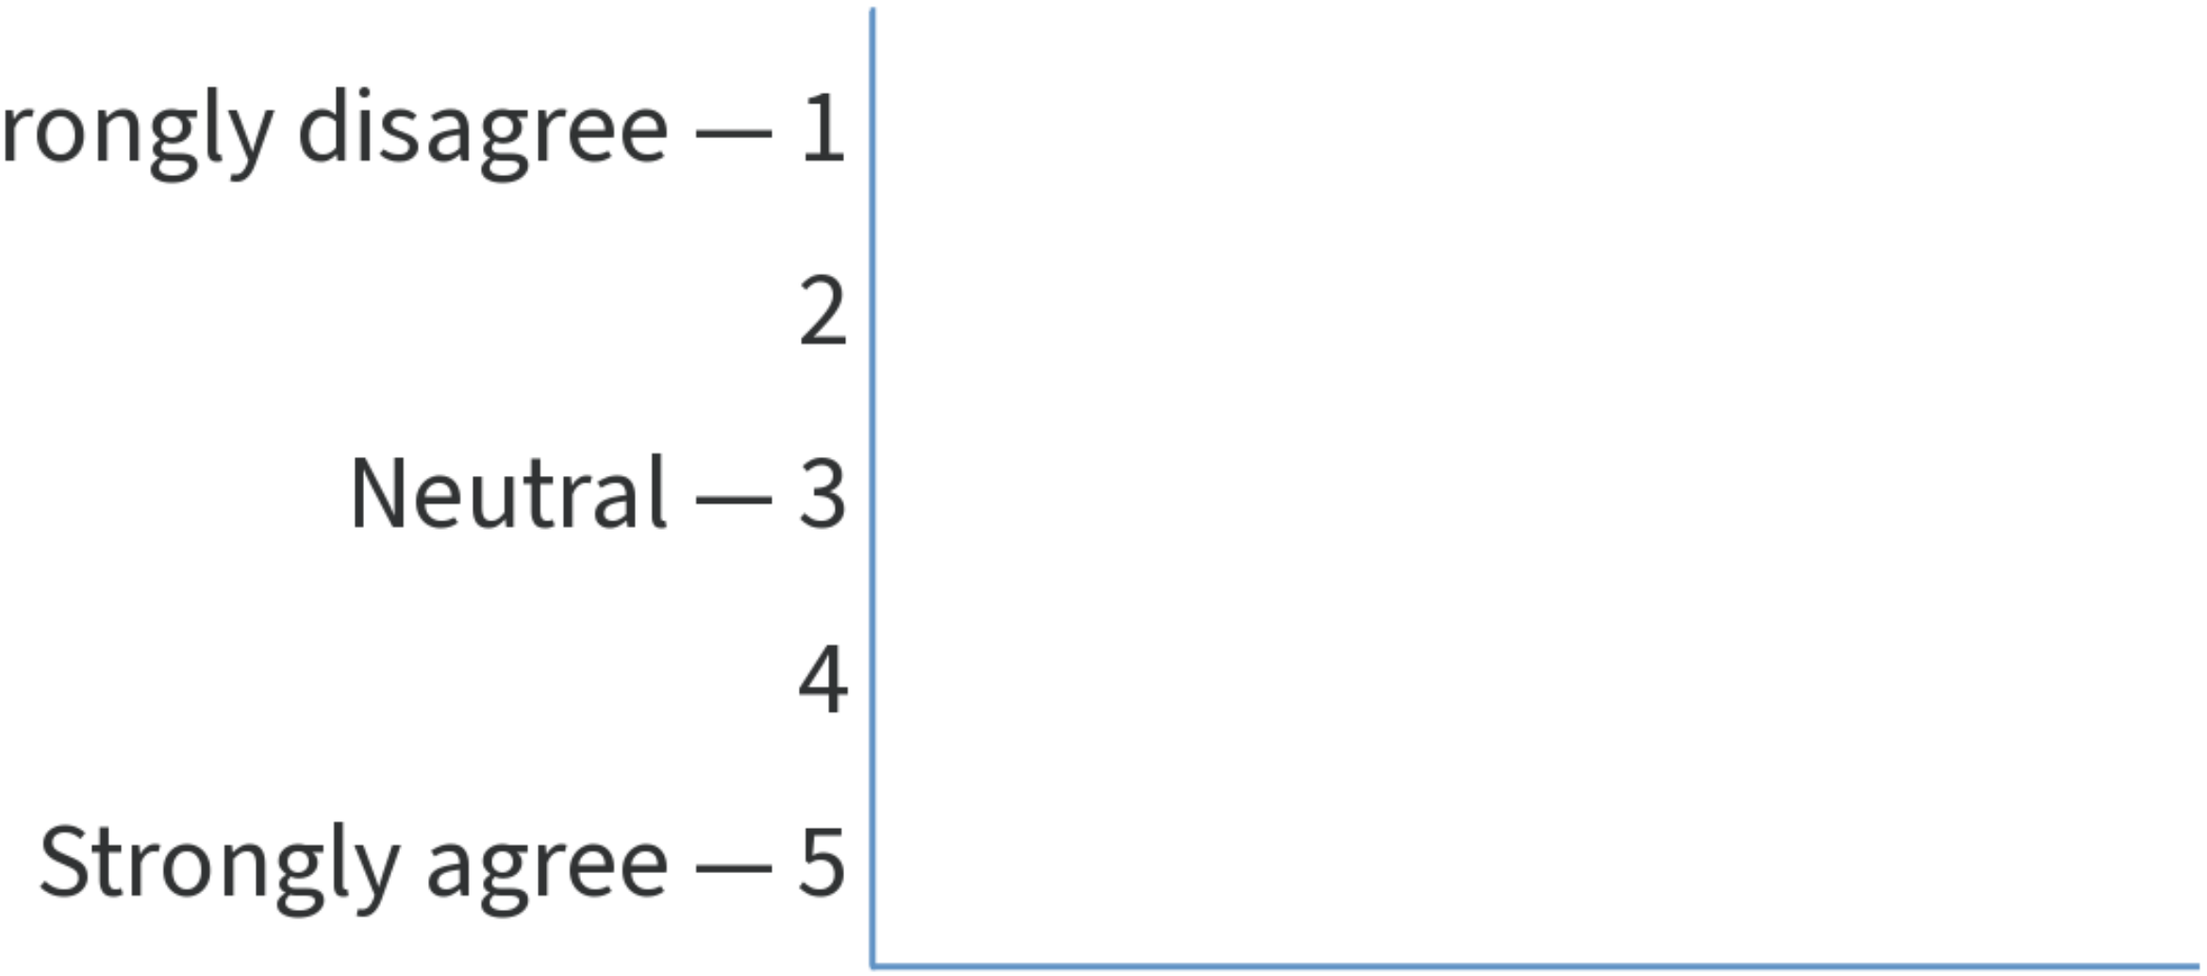

Total Results: 0

ChatGPT understands English

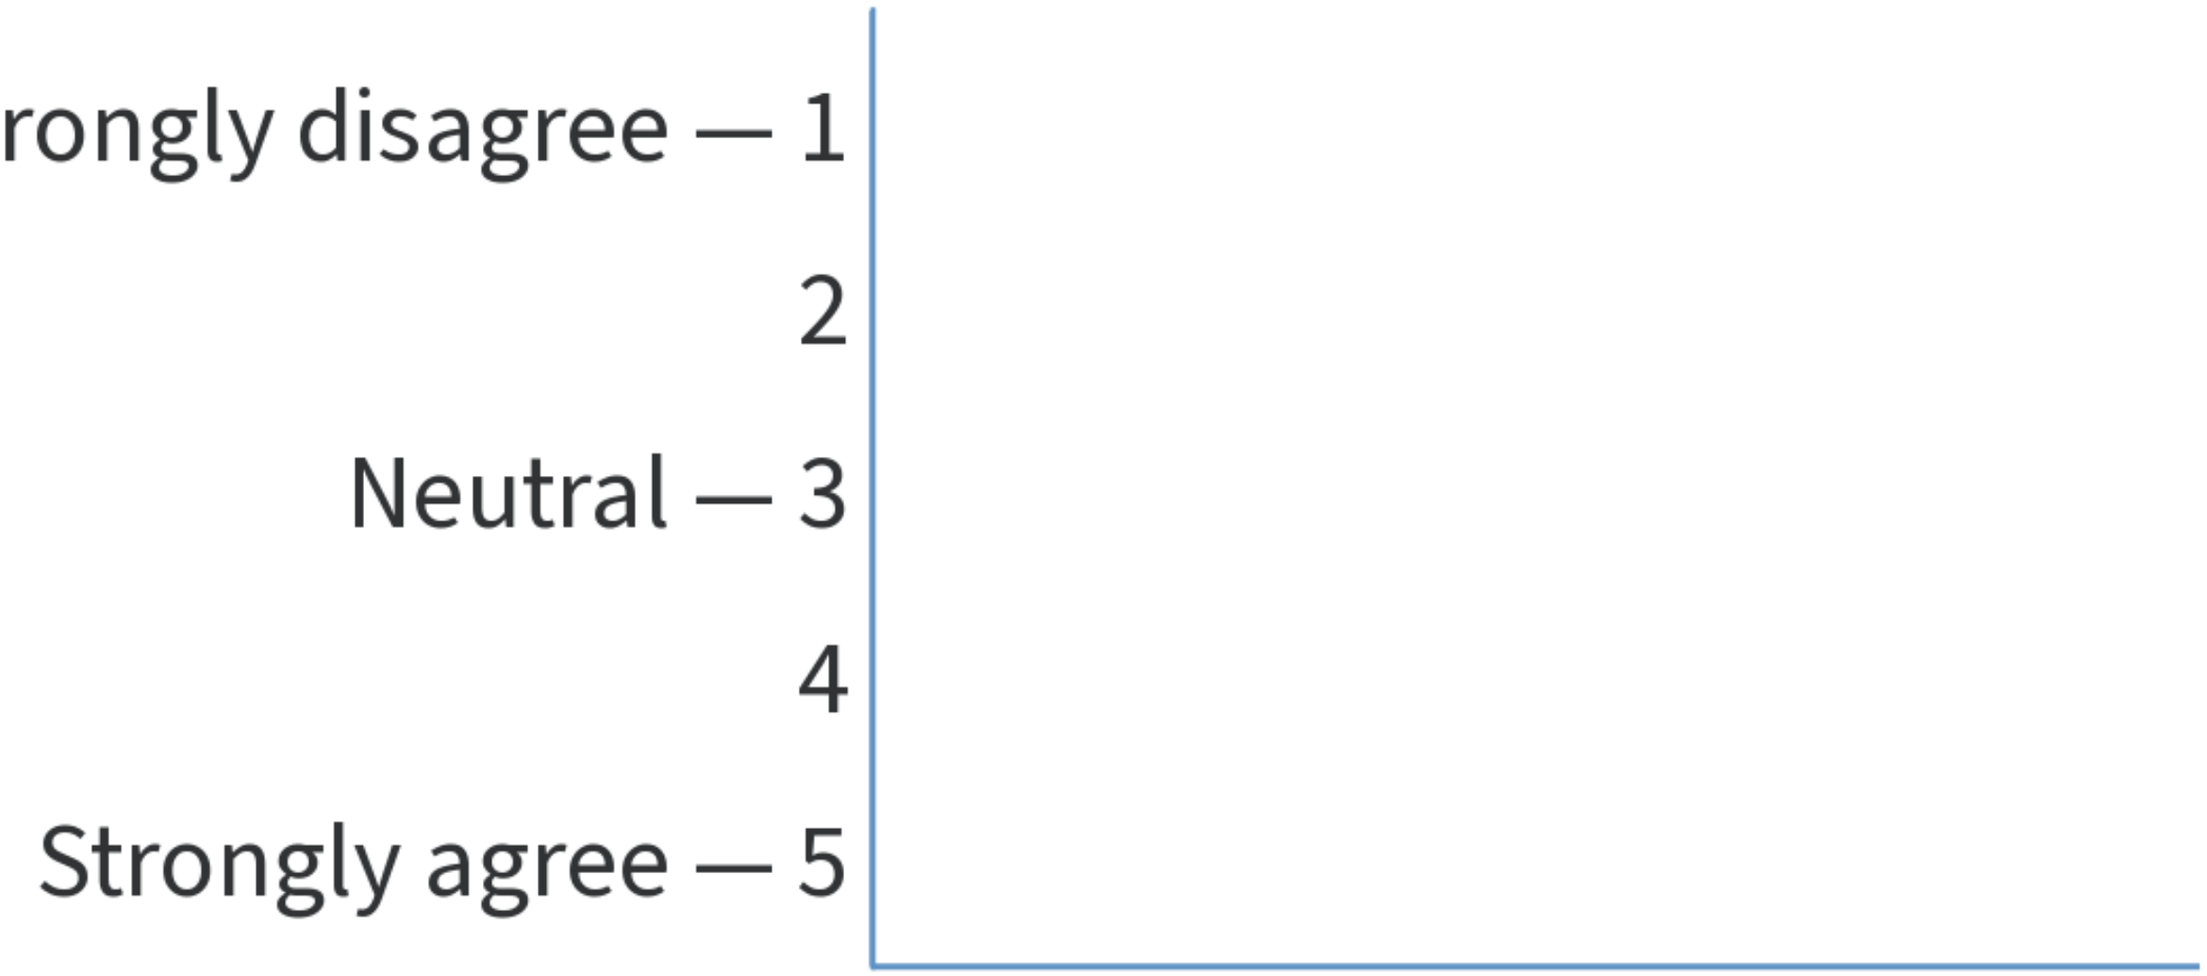

Total Results: 0

You can trust what ChatGPT says

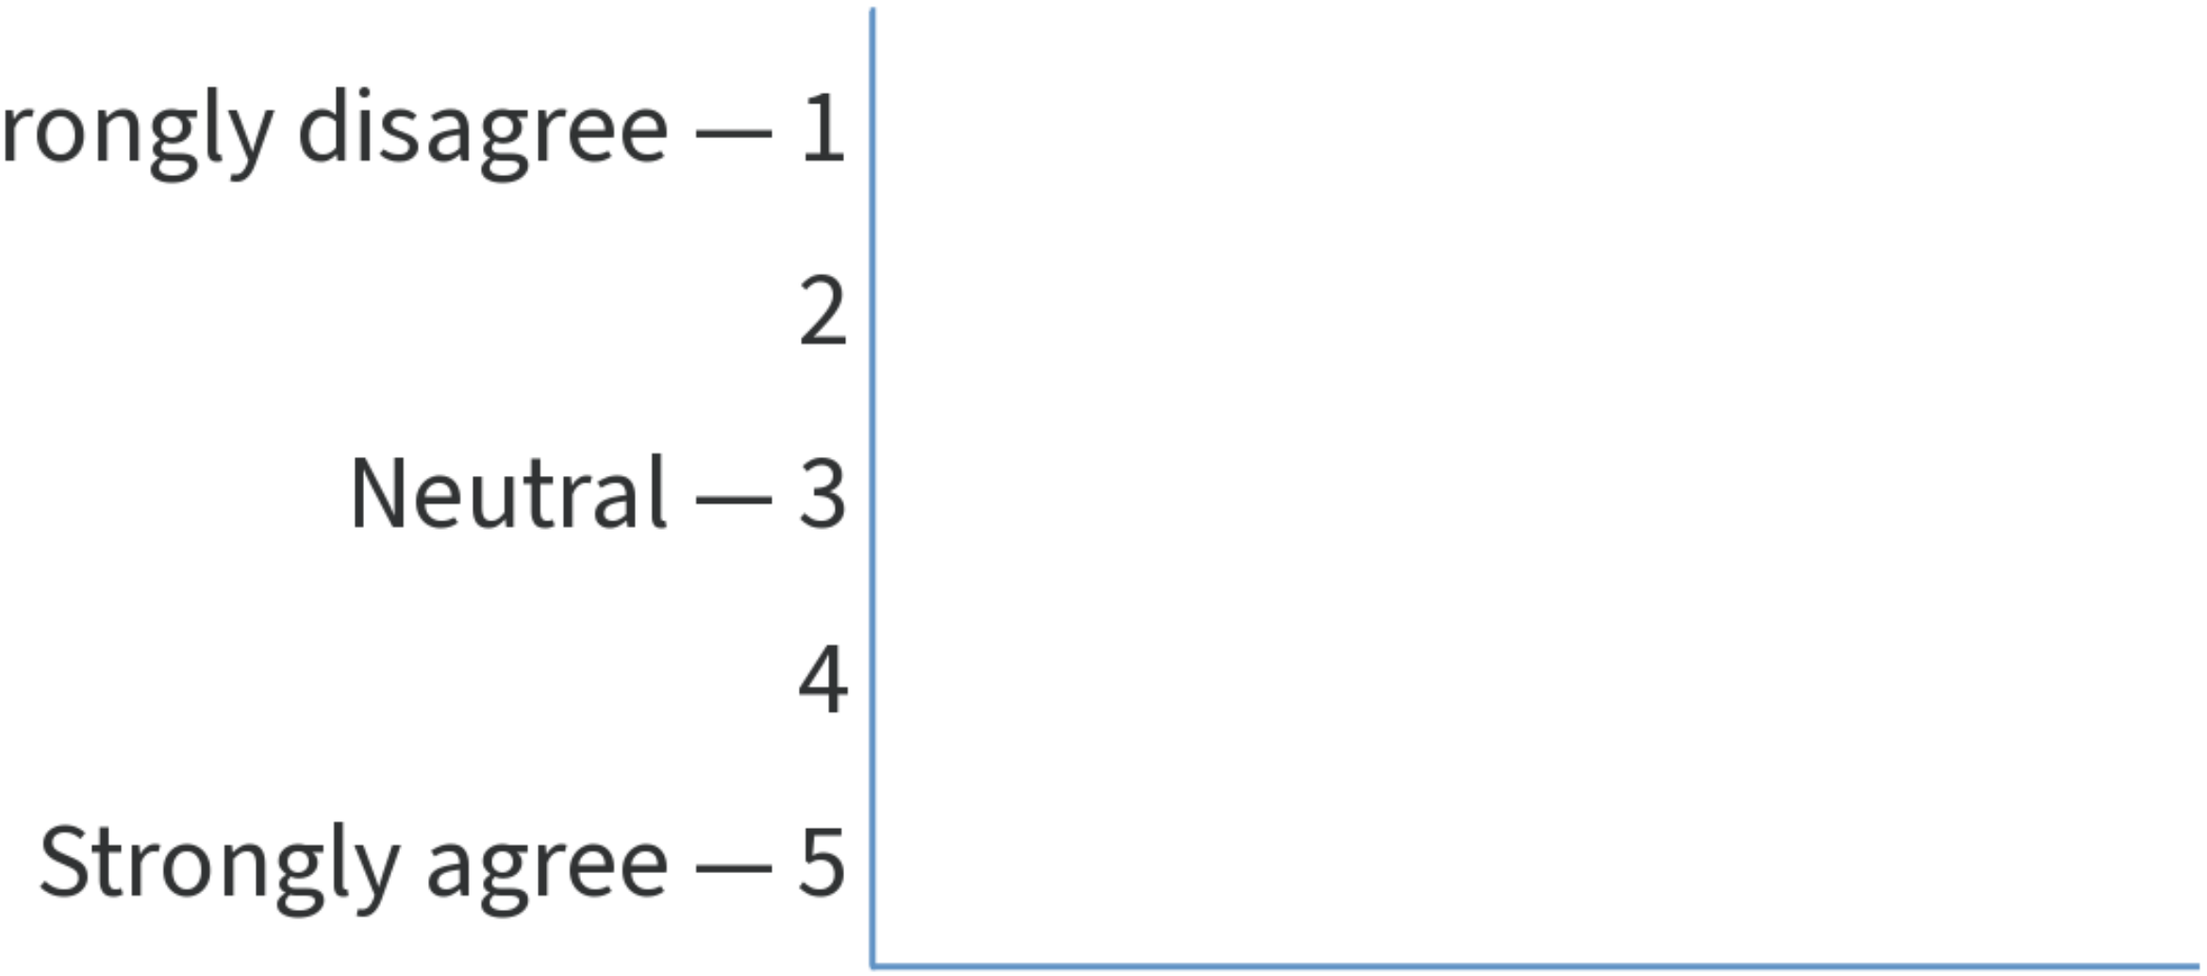

Total Results: 0

What is ChatGPT good for?

Total Results: 0

What are the risks associated with ChatGPT?

Total Results: 0

Overview

- What is Computational Linguistics
- Who’s here
- Syllabus
- Showing the computer who’s boss

What is Computational Linguistics?

- Getting computers to deal with human languages
  - ... for practical applications (examples?)
  - ... for linguistic research (examples?)

Linguistic research

- Searching large corpora for patterns of use and linguistic examples
- Creating structured databases of information for typological research (Autotyp, ODIN)
- Creating ontologies for interoperable markup of linguistic resources (GOLD)
- Modeling human linguistic competence and performance (computational psycholinguistics, grammar engineering)
- Software to facilitate language documentation (Elan, FIELD, SIL FieldWorks, Grammar Matrix, AGGREGATION, EL-STEC)

Practical applications

- Speech recognition
- Speech synthesis
- Machine translation
- Information retrieval
- Natural language interfaces to computers
- Dialogue systems

Practical applications

- Computer-assisted language learning (CALL)
- Grammar checkers
- Spell checkers
- OCR (optical character recognition)
- Handwriting recognition
- Augmentative and assistive communication

Practical applications

- BioMedical NLP: Matching patients to clinical trials
- BioMedical NLP: Flagging electronic health records for urgent tests
- BioMedical NLP: Assistance in coding for insurance billing
- BioMedical NLP: Searching the biomedical literature for untested but promising things to study
- Legal domain: Electronic discovery

Practical applications

- B2B: Sentiment analysis for brand tracking
- Context-aware advertising
- Intelligence/national security: Monitoring social media, news, intercepted email/voice traffic
- ...

End-to-end applications can be constructed from components that handle subtasks

- Each subtask has input and output
- Each subtask can be evaluated
    - often: precision, recall
    - intrinsic and extrinsic evaluation
- Output from one subtask is input to the next (in pipeline models)
- Many subtasks have “analysis” and “generation” variants
- Examples of subtasks?

Subtasks
(What's the input? What's the output?)

- Part of Speech tagging
- Named Entity Recognition
- Lemmatization
- Word sense disambiguation
- Event detection
- Dialog act labeling

- Morphological analysis
- Parsing (constituent structure, dependency structure)
- Coreference resolution
- Language modeling
- Alignment (of bitexts)
- ...

Statistical v. symbolic methods

- Statistical methods involve *training a stochastic model* on a body of data so it can predict the most probable label/structure/etc for new data
    - Knowledge comes from implicit patterns in naturally occurring language (unsupervised learning) or from hand-labeled data (supervised learning)
- Symbolic methods involve *knowledge engineering*, or hand-coding of linguistic knowledge which is then applied to tasks
- Statistical methods provide *robustness*, symbolic methods *precision*
- Statistical and symbolic methods can be combined

The World of CL: CL at UW

- Linguistics (CLMS)
- CSE
- ECE
- Biomedical informatics
- iSchool

The World of CL: CL in Seattle

- Microsoft (MSR)
- Amazon
- AI2
- Facebook/Meta
- Google
- ...

World of CL: ACL

- Association for Computational Linguistics; our chapter: NAACL
- Conferences: ACL, NAACL, EACL, IJCNLP, EMNLP, others
- Workshops
- Publications
  - *Computational Linguistics* and *TACL* journals
  - Conference and workshop proceedings: ACL Anthology https://www.aclweb.org/anthology/

World of CL: online communication

- @UW: cl-announce http://mailman.u.washington.edu/mailman/listinfo/cl-announce

- International: corpora https://list.elra.info/mailman3/postorius/lists/corpora.list.elra.info/

- Twitter: #NLProc, conference hashtags

Learning Outcomes

- Become familiar with topics and research methods in computational linguistics and related fields
- Be able to conceptualize research problems from the perspective of computational linguistics
- Understand ways to computationally encode linguistic knowledge
- Practice designing and carrying out linguistically-informed analyses of NLP systems
- Practice identifying and scoping personal interests in the field
- Be an informed consumer of NLP/speech technology and popular press reporting on NLP/speech technology

Why this class is weird

- Upper-division survey course
- Students with diverse backgrounds
- So why teach this as one cross-listed course?

Syllabus

- Web page: http://courses.washington.edu/ling472
    - NB: Things are due already this week (RQ, Assignment 0)
- Slides will be posted (often before lecture)
- Using Canvas and Zoom (recordings will be included on Canvas page)
- Lab meetings (Fridays)

Course requirements

- Homework assignments (5 total, turned in via Canvas): 50%
  - Coding *and* writing: Writing will be 50% of the grade
- Final project: 30%
- Reading questions: 5%
- Blog assignment: 5%
- Beginning & end quarter reflections: 10%
- Up to 2% adjustment for:
  - Extra credit points for original clarification questions
  - In class participation
    - Other on-line participation
- Get set up: see course web page for server cluster accounts, lab access, reading assignments, etc.

Course requirements

- Homework assignments (5 total, turned in via Canvas): 60%
  - Coding *and* writing: Writing will be 50% of the grade
- Final project: 30%
- Reading questions: 5%
- Blog assignment: 5%
- Up to 2% adjustment for:
  - Extra credit points for original clarification questions
  - In class participation
    - Other on-line participation
- Get set up: see course web page for server cluster accounts, lab access, reading assignments, etc.

To course web page

- Reading questions
- Blog assignment
- Reflections assignment
- Term project

Letting the computer know who's boss

- Computer 'literacy' is really a combination of experience and attitude

- Experience gives you the answers to many questions and a sense of what the possible space of answers is

- The important attitude boils down to confidence in one's ability to find the answer to a new question

- There are always new questions because:

  - The technology is always developing

  - There is too much for any one person to know it all

Letting the computer know who's boss

- Keep in mind:
  - It's always obvious once you know the answer
  - All pieces of software were designed by some person or people with some functionality in mind
- Places to look for answers:
  - On-line documentation (man, info, help)
  - Product websites (esp. discussion forums)
  - Google: websites, and especially newsgroups
  - Off-line documentation (i.e., books!)
  - Work together!
    - ... and post to the discussion boards in Canvas
  - 10 minute rule
    - It's ~~okay~~ critically important to ask questions!

Questions?

Overview

- What is Computational Linguistics
- Syllabus
- Who’s here
- Showing the computer who’s boss