More automated analyses

Hakyung Sung
University of Hawaii at Manoa

November 18, 2026

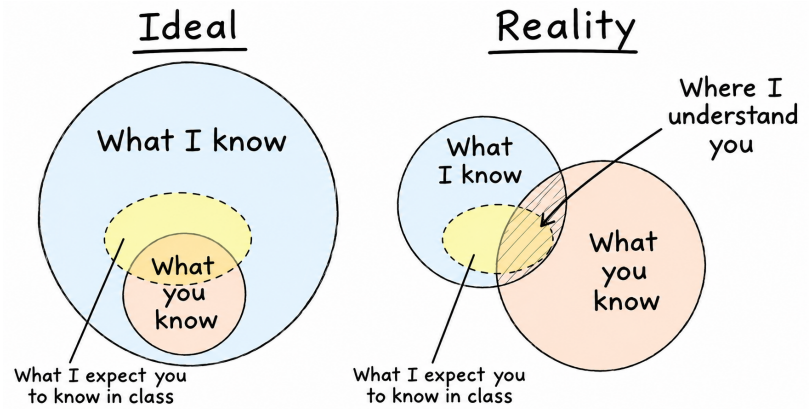
Topics to cover

- Projects (10 mins)
- Review: Semantic role labeling (5 mins)
- Sentiment, discourse analysis (15 mins)
- Presentations (40 mins)
- **Lab (most time)**

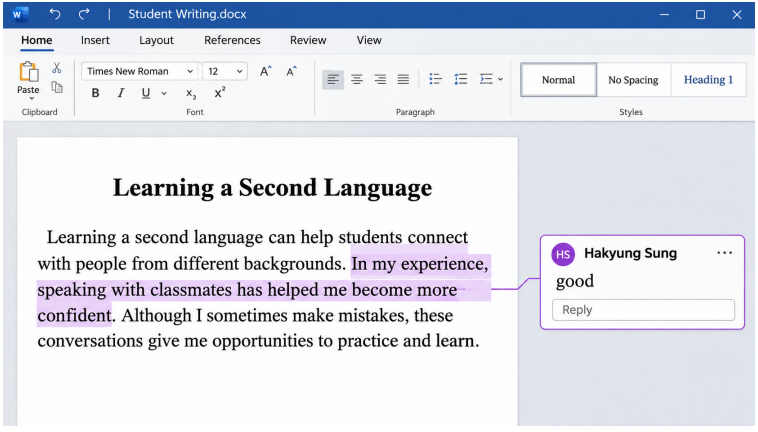
Projects

- Overall, you’re making great progress!
- Recommendation: Continue your current topic unless you have a VERY strong reason to change it.
- Hope my comments are helpful (but they may not always).

I only know what I know.



*__Good__ = I understood this clearly.



Writing a paper is an act of communication, beyond NLP approaches and numbers from data.

What can you do (some corny advice)?

- Read (with attention to how authors explain their ideas).
- Write (with your reader(s) in mind).
- Revise many times.

- Ask me questions when you need guidance (*this also helps me understand what you know and/or want to know).

Working toward final report

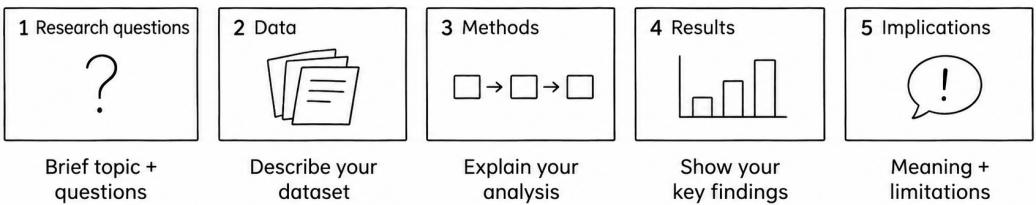
- Final presentation (12/9)
- Final report (12/14)
- Descriptions on course website: here and here

Final presentation

- Each student will present their project to the class.
- Each 10-12 mins.
- Due to time constraints, audit student are exempted. So we have 12 in total.

Final presentation

Presentation: 5 slides



Final presentation

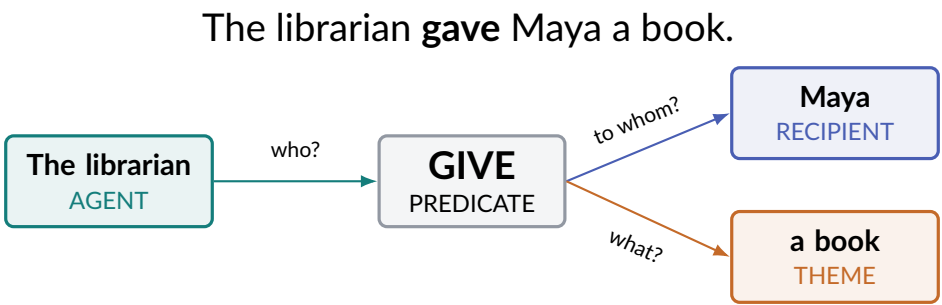
- Complete the google form (just like you did for the midway presentation).
- Please submit the slides by **the day before class (12/8)**.

Final report

- Submit a complete report (the most important assignment for this course!) with all sections needed for an academic paper.
- Share your analysis code (written in Python) and a README in a GitHub repository.
  - The code should correspond to the procedures in your Methods section.
  - The README should explain what the code does and how to run it.
  - We will cover GitHub and README files in Week 15 (December 2). For now, keep your code organized.
- Acknowledge any AI tools you used and explain how you used them (e.g., coding, revising your writing, generating workflow figure).
- Optional: If you have a particular audience, journal, or conference in mind for your paper, please let me know.
- Any questions?

Review: Semantic role labeling

Semantic roles: *Who did what to whom?*



Semantic roles describe participants' relationships to an event.

Other semantic roles

- **Force** (non-volitional cause)
  The wind opened the door.

- **Agent** (intentional actor)
  Maya opened the door.

- **Instrument** (tool used)
  Maya opened the door with a key.

- **Beneficiary** (who benefits)
  Ann made reservations for her boss.

- **Source** and **Goal** (movement)
  I flew from Boston to Portland.

Different syntax, Same semantic roles

- **Sentence A: Double-object construction**

  The librarian gave Maya a book.

- **Sentence B: Prepositional dative** (dative alternations)

  The librarian gave a book to Maya.

Agent **Theme** **Recipient**

Semantic role annotation

- *PropBank* annotates predicates and their arguments with semantic roles. Each predicate sense has its own roleset (a set of argument definitions).

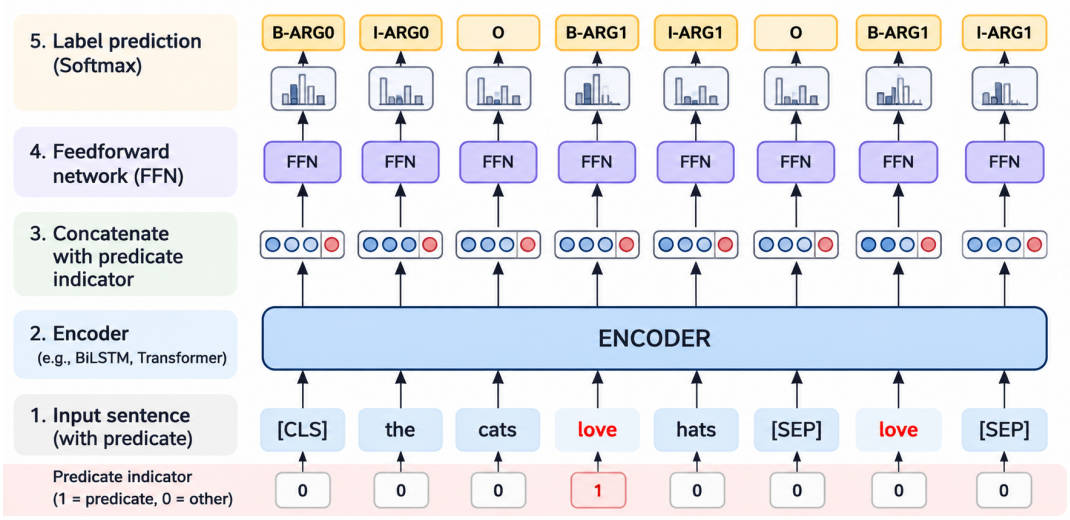
- **Core arguments** are labeled using numbers:
  - ARG0: Typically Proto-Agent
  - ARG1: Typically Proto-Patient
  - ARG2–ARG4: Other predicate-specific roles

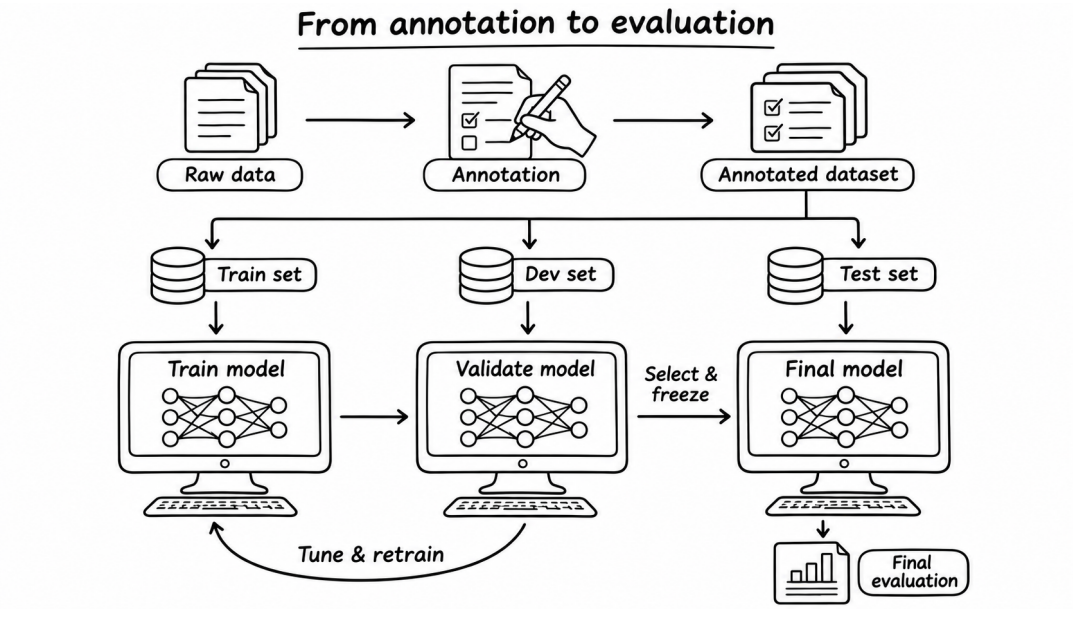
- **Modifier arguments (ARGM)**
  - **ARGM-TMP**: Time (when?)
  - **ARGM-LOC**: Location (where?)
  - **ARGM-DIR**: Direction (where to/from?)
  - **ARGM-MNR**: Manner (how?)
  - **ARGM-CAU**: Cause (why?)
  - **ARGM-PRP**: Purpose (for what purpose?)

Palmer, M., Gildea, D., & Kingsbury, P. (2005). The Proposition Bank: An annotated corpus of semantic roles. Computational Linguistics, 31(1), 71–106.

Semantic role labeling (SRL)

- Automatically identify predicates and assign semantic roles to their arguments.
- Traditionally trained on annotated resources (e.g., PropBank).


5. Label prediction (Softmax)
4. Feedforward network (FFN)
3. Concatenate with predicate indicator
2. Encoder (e.g., BiLSTM, Transformer)
1. Input sentence (with predicate)
Predicate indicator (1 = predicate, 0 = other)
B-ARG0
I-ARG0
O
B-ARG1
I-ARG1
O
B-ARG1
I-ARG1
FFN
ENCODER
[CLS]
the
cats
love
hats
[SEP]
love
[SEP]
0
0
0
1
0
0
0
0


From annotation to evaluation
Raw data
Annotation
Annotated dataset
Train set
Dev set
Test set
Train model
Validate model
Select & freeze
Final model
Tune & retrain
Final evaluation

Sentiment, discourse analysis

Sentiment analysis

- Language conveys not only information but also feelings and attitudes.
- This field is related to affective computing.
- Simple Q: Is the writer/speaker *positive* or *negative* about something?
- Related concepts:
  - *Emotion*: a response to a specific event (e.g., anger, joy).
  - *Mood*: a more diffuse state (e.g., cheerful, gloomy).
  - *Interpersonal stance*: an orientation toward others (e.g., supportive, distant).
- Connotation: words can carry evaluative associations beyond their basic meanings.

Sentiment depends on the target

e.g.,

I love the camera, but its battery is terrible.

- Document-level analysis: assigns an overall sentiment to the text.
- Aspect-level analysis: distinguishes evaluations of different targets.
  - Camera $\rightarrow$ positive (*love*).
  - Battery $\rightarrow$ negative (*terrible*).

Two approaches

- Supervised machine learning
  - Train a classifier on *texts labeled for sentiment*.
  - Learn patterns associated with positive and negative texts.
  - Apply the model to new texts.

- Lexicon-based analysis
  - Look up *words in a sentiment* or affect dictionary.
  - Combine word-level scores to characterize a text.

How are sentiment lexicons built?

- Human annotation: collect judgments of words' affective meanings.
  - Best–worst scaling: select the most and least positive words from a small set.
- Bootstrapping: expand *seed* words using relationships such as synonyms and antonyms (from word embeddings).
- Supervised induction: learn word associations from labeled texts, such as rated reviews.
- Corpus sourced from web: `https://hedonometer.org/timeseries/en_all/?from=2021-12-30&to=2023-06-29`

Applications?

Q. What questions in L2 research could we explore using sentiment analysis?

(Ha will lead the discussion after the presentation)

Presentation 1: Taboada (2016)

- A review of how linguistic knowledge contributes to sentiment analysis.

- e.g., Stewart (2015) analyzes students' written comments in course evaluations using the Appraisal framework (Martin & White 2005).

Table 1: Resources of attitude and graduation.

Domains of APPRAISAL	Sub-domains of APPRAISAL	Resources	Examples
Attitude	Affect	Happy, sad, love, hate, long for, fearful, confident, trusting, anxious, startled, satisfied, involved, bored with, flat, etc.	She's **happy**. They were **bored with** action films.
	Judgement	Lucky, unlucky, powerful, weak, careful, reckless, honest, dishonest, good, bad, etc.	The men were **honest**. They're **bad** bankers.
	Appreciation	Exciting, boring, lovely, ugly, balanced, unbalanced, clear, unclear, effective, ineffective, etc.	The football match is **exciting**. His approaches were **effective**.
Graduation	Force	*Very*, *somewhat*, *slightly*, *greatly*, *more*, *less*, etc.	This upset me *greatly*. She's *very* happy.
	Focus	*True*, *real*, *of sorts*, etc.	He's a *true* friend. It was an apology *of sorts*.

Table from Chung & Crosthwaite (2024). Stance-taking through APPRAISAL in L1 and L2 English argumentative essays

Discourse analysis

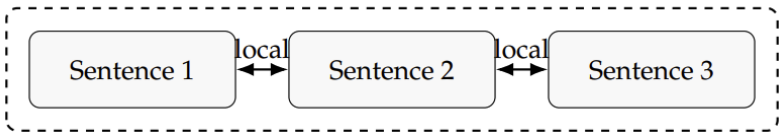
How do ideas connect?

e.g.,

Jane took a train to Istanbul. She likes spinach.

Jane took a train to Istanbul. She had to attend a conference.

Local coherence? Global coherence?



Global coherence concerns how the full sequence supports a larger purpose or structure

Figure 11: Local vs. global coherence

Ways to represent discourse relations

- Rhetorical Structure Theory (RST)
  - Relates text spans within a hierarchical structure.
  - Often distinguishes a central message (*nucleus*) from supporting material (*satellite*).

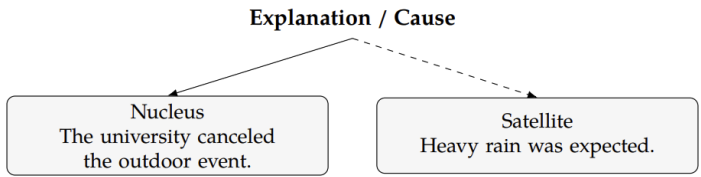
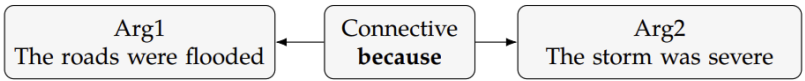


Figure 12: A simplified nucleus-satellite illustration of an explanatory relation.

Ways to represent discourse relations

- Penn Discourse Treebank (PDTB)
  - Identifies relations between two argument spans.
  - Often uses connectives (lexically grounded) such as *because*, *although*, and *when*.



Relation: causal explanation

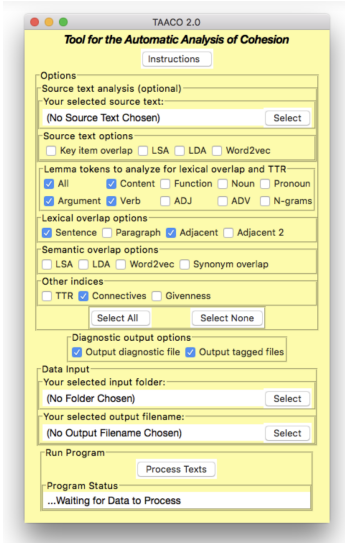
Figure 13: Schematic PDTB-style representation of two argument spans and a connective.

NLP tool to evaluate coherence

TAACO: TOOL FOR THE AUTOMATIC ANALYSIS OF COHESION

TAACO is an easy to use tool that calculates 150 indices of both local and global cohesion, including a number of type-token ratio indices (including specific parts of speech, lemmas, bigrams, trigrams and more), adjacent overlap indices (at both the sentence and paragraph level), and connectives indices.

TAACO takes plain text files as input (it will process all plain text files in a particular folder) and produces a comma separated values (.csv) spreadsheet that is easily read by any spreadsheet software.



Automatic discourse parsing and evaluation

- An RST-oriented parser typically:
  1. Segments text into *elementary discourse units* (EDUs).
  2. Predicts relations and structures linking those units.

- EDUs often resemble clauses, but boundaries depend on the annotation scheme.

- Evaluation:
  - Compare predicted boundaries and relations with human annotations.
  - Compare coherence scores with human ratings.
  - Test whether a model favors original sentence order over shuffled versions.

Global coherence: Argument mining

- Identify claims, reasons, and support or attack relations.
  - Claim: The program should be expanded.
  - Support: Participation has increased.
  - Support: Participants report benefits.

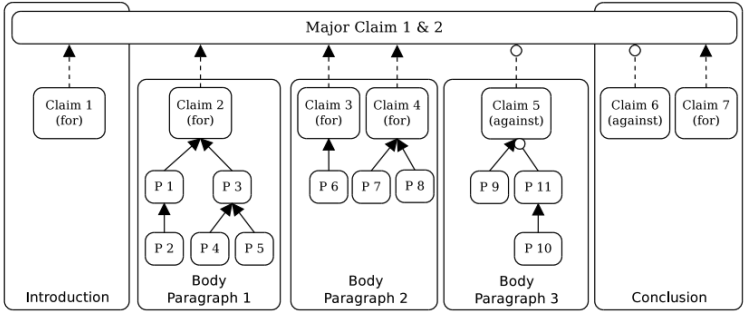


Figure 2
Argumentation structure of the example essay. Arrows indicate argumentative relations. Arrowheads denote argumentative support relations and circleheads attack relations. Dashed lines indicate relations that are encoded in the stance attributes of claims. "P" denotes premises.

Global coherence: Scientific report

Scientific Discourse Tagging for Evidence Extraction

Xiangci Li[1]* **Gully Burns**[2] **Nanyun Peng**[3]
[1] University of Texas at Dallas, [2] Chan Zuckerburg Initiative
[3] University of California Los Angeles
lixiangci8@gmail.com,gully.burns@chanzuckerburg.com
violetpeng@cs.ucla.edu

lc3 , the mammalian atg8 homolog , undergoes a set of modifications resulting in conversion from lc3i to lc3ii during autophagy 42 . **[fact]** to further test the function of rag in autophagy **[goal]** we examined the lc3 modification in hek293 cells . **[method]** expression of raga ql and ragc sn inhibited lc3 conversion in response to amino acid starvation (fig. 7e) . **[result]** furthermore , expression of raga tn and ragc ql enhanced lc3 conversion even in the presence of amino acids . **[result]** these results are consistent with the data observed in drosophila and further demonstrate a role of the rag gtpases in autophagy regulation in response to nutrient signals **[implication]**

Figure 1: An example paragraph tagged with scientific discourse tags on each clause in SciDT dataset (Dasigi et al., 2017). The text is tokenized and converted to lower case.

Figure 1 from Li, Burns & Peng (2021), Scientific Discourse Tagging for Evidence Extraction

Applications?

Q. What questions in L2 research could we explore using discourse analysis?

Presentation 2: Fischer & Biemann (2025)

- Help researchers annotate sentences using the Discourse Analysis Tool Suite (DATS).
- Good source for available datasets.
- More information will on Jue's presentation