Using Sentence-Level LSTM Language Models for Script Inference

Karl Pichotta and Raymond J. Mooney
The University of Texas at Austin

ACL 2016, Berlin

Event Inference: Motivation

- Suppose we want to build a Question Answering system…

Event Inference: Motivation

- *The Convention ordered the arrest of Robespierre.… Troops from the Commune, under General Coffinhal, arrived to free the prisoners and then marched against the Convention itself.*

  –Wikipedia

- **Was Robespierre arrested?**

Event Inference: Motivation

- *The Convention **ordered the arrest** of Robespierre.… Troops from the Commune, under General Coffinhal, arrived to free the prisoners and then marched against the Convention itself.*

  –Wikipedia

- **Was Robespierre arrested?**

Event Inference: Motivation

- *The Convention **ordered the arrest** of Robespierre.… Troops from the Commune, under General Coffinhal, **arrived to free the prisoners** and then marched against the Convention itself.*

  –Wikipedia

- **Was Robespierre arrested?**

Event Inference: Motivation

- *The Convention **ordered the arrest** of Robespierre.… Troops from the Commune, under General Coffinhal, **arrived to free the prisoners** and then marched against the Convention itself.*

  –Wikipedia

- **Was Robespierre arrested? Very probably!**

Event Inference: Motivation

- *The Convention **ordered the arrest** of Robespierre.… Troops from the Commune, under General Coffinhal, **arrived to free the prisoners** and then marched against the Convention itself.*

  –Wikipedia

- **Was Robespierre arrested? Very probably!**
  - **…But this needs to be inferred.**

Event Inference: Motivation

- Question answering requires inference of probable implicit events.

- We'll investigate such event inference systems.

Outline

- Background & Methods
- Experiments
- Conclusions

Outline

- Background & Methods
- Experiments
- Conclusions

Outline

- Background & Methods
- Event Sequence Learning & Inference
- Sentence-Level Language Models

Outline

- Background & Methods
- Event Sequence Learning & Inference
- Sentence-Level Language Models

Event Sequence Learning

- [Schank & Abelson 1977] gave a non-statistical account of *scripts* (events in sequence).

- [Chambers & Jurafsky (ACL 2008)] provided a statistical model of *(verb, dependency)* events.

- A recent body of work focuses on learning statistical models of event sequences [e.g. P. & Mooney (AAAI 2016)].

- *Events* are, for us, verbs with multiple NP arguments.

Event Sequence Learning

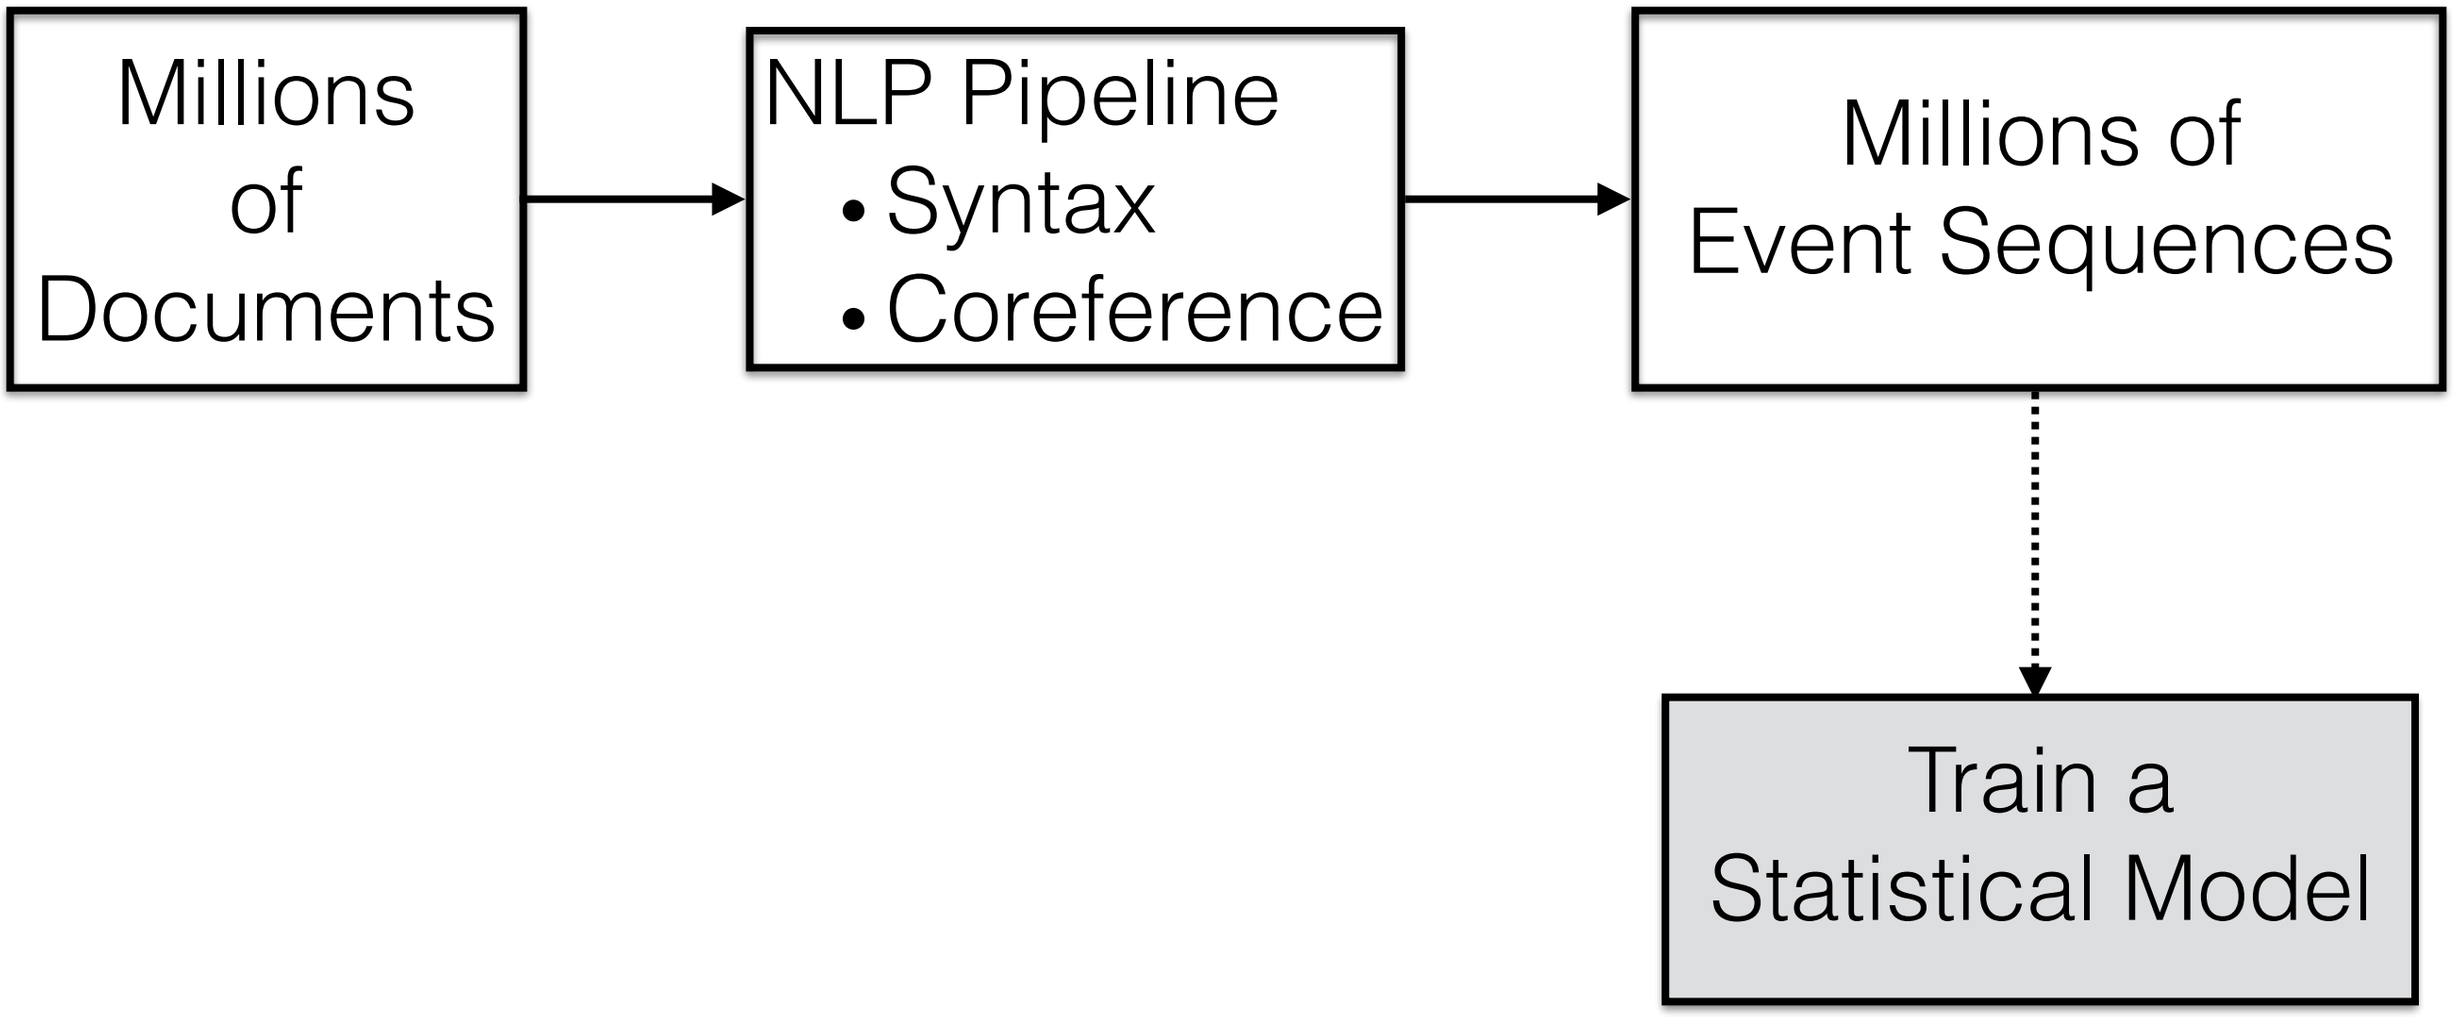

Event Sequence Inference

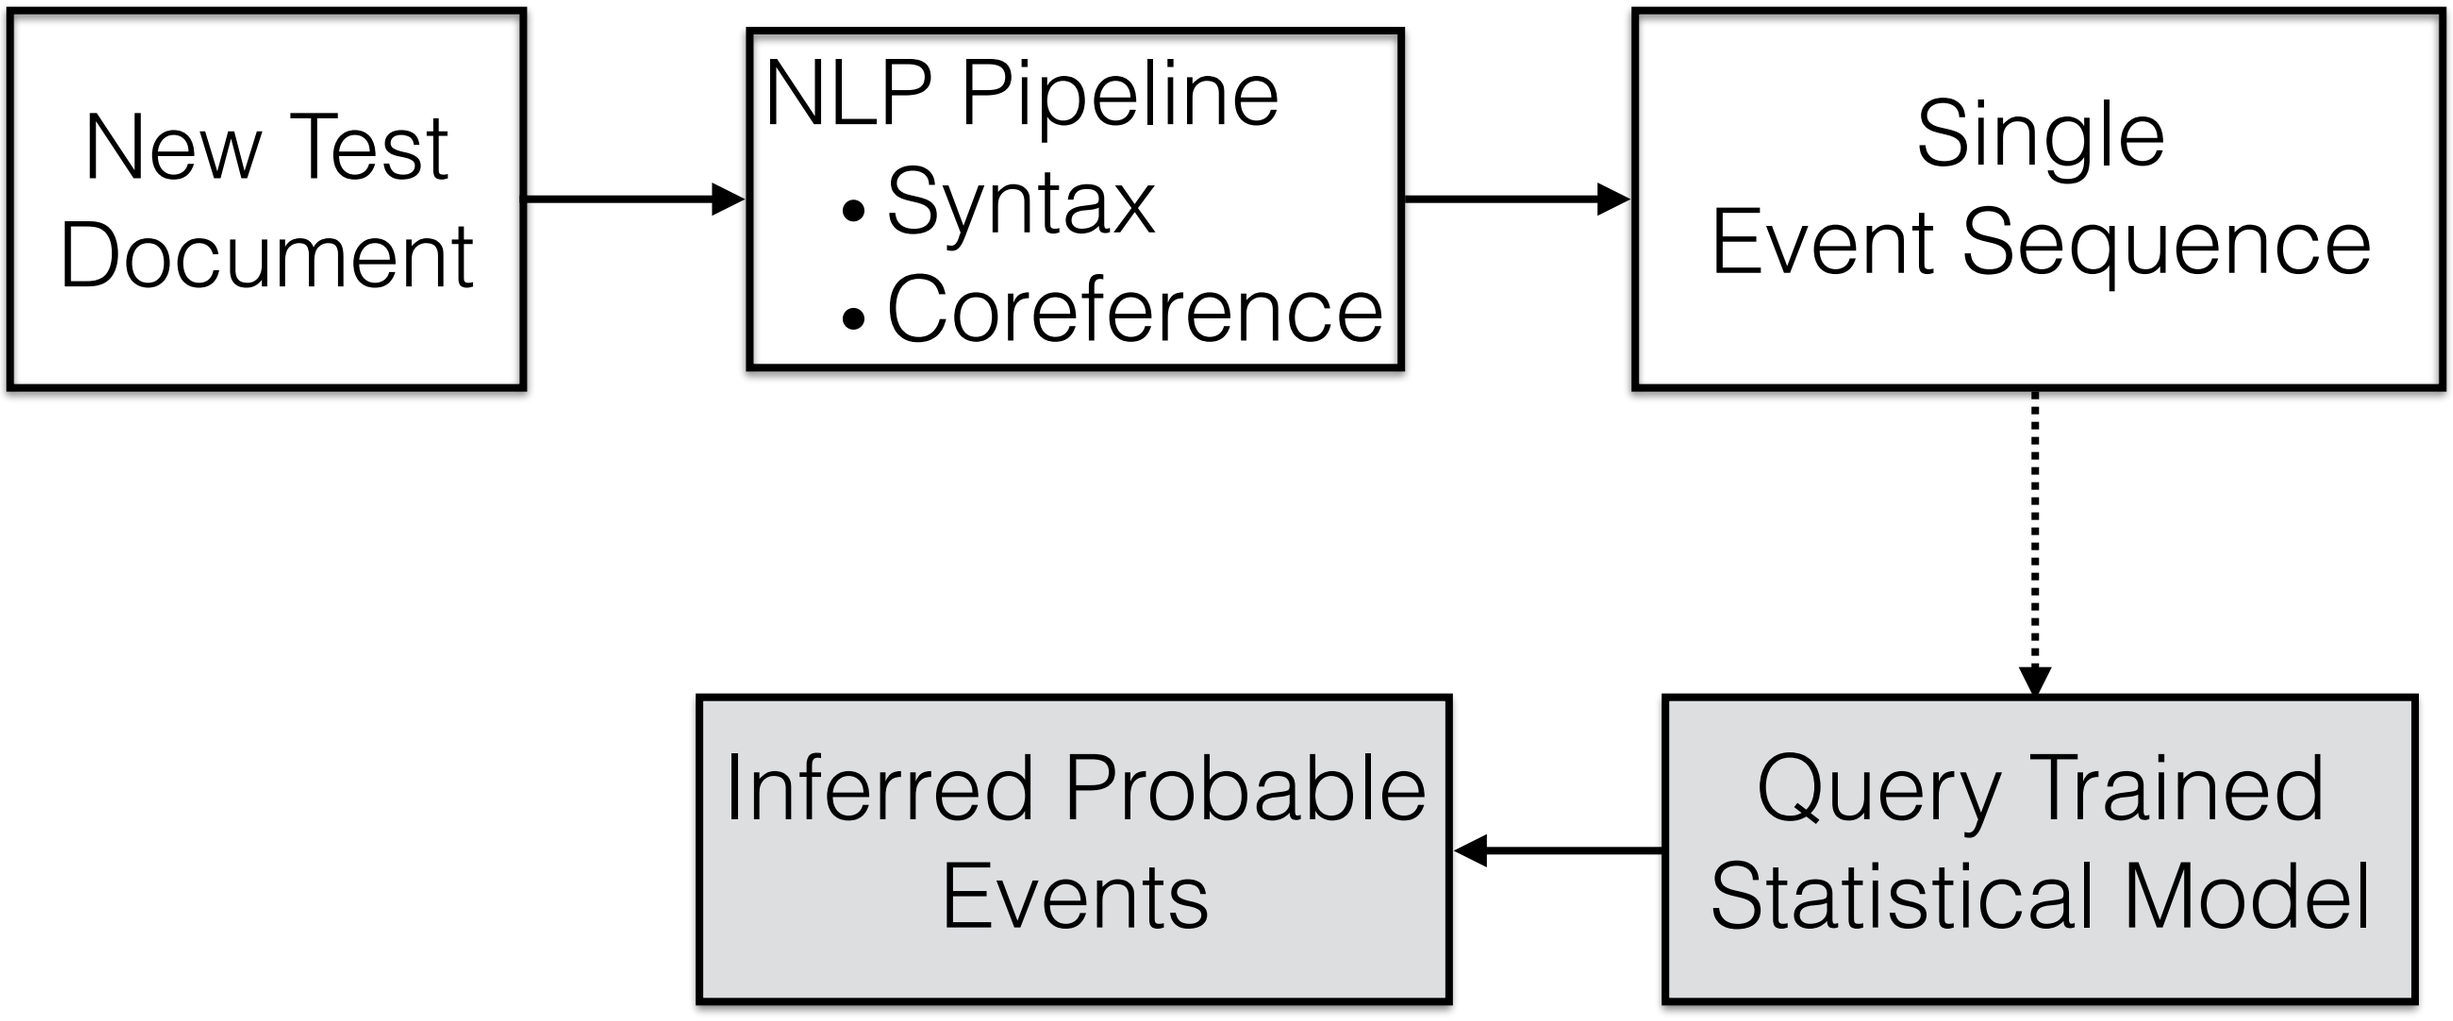

Event Sequence Inference

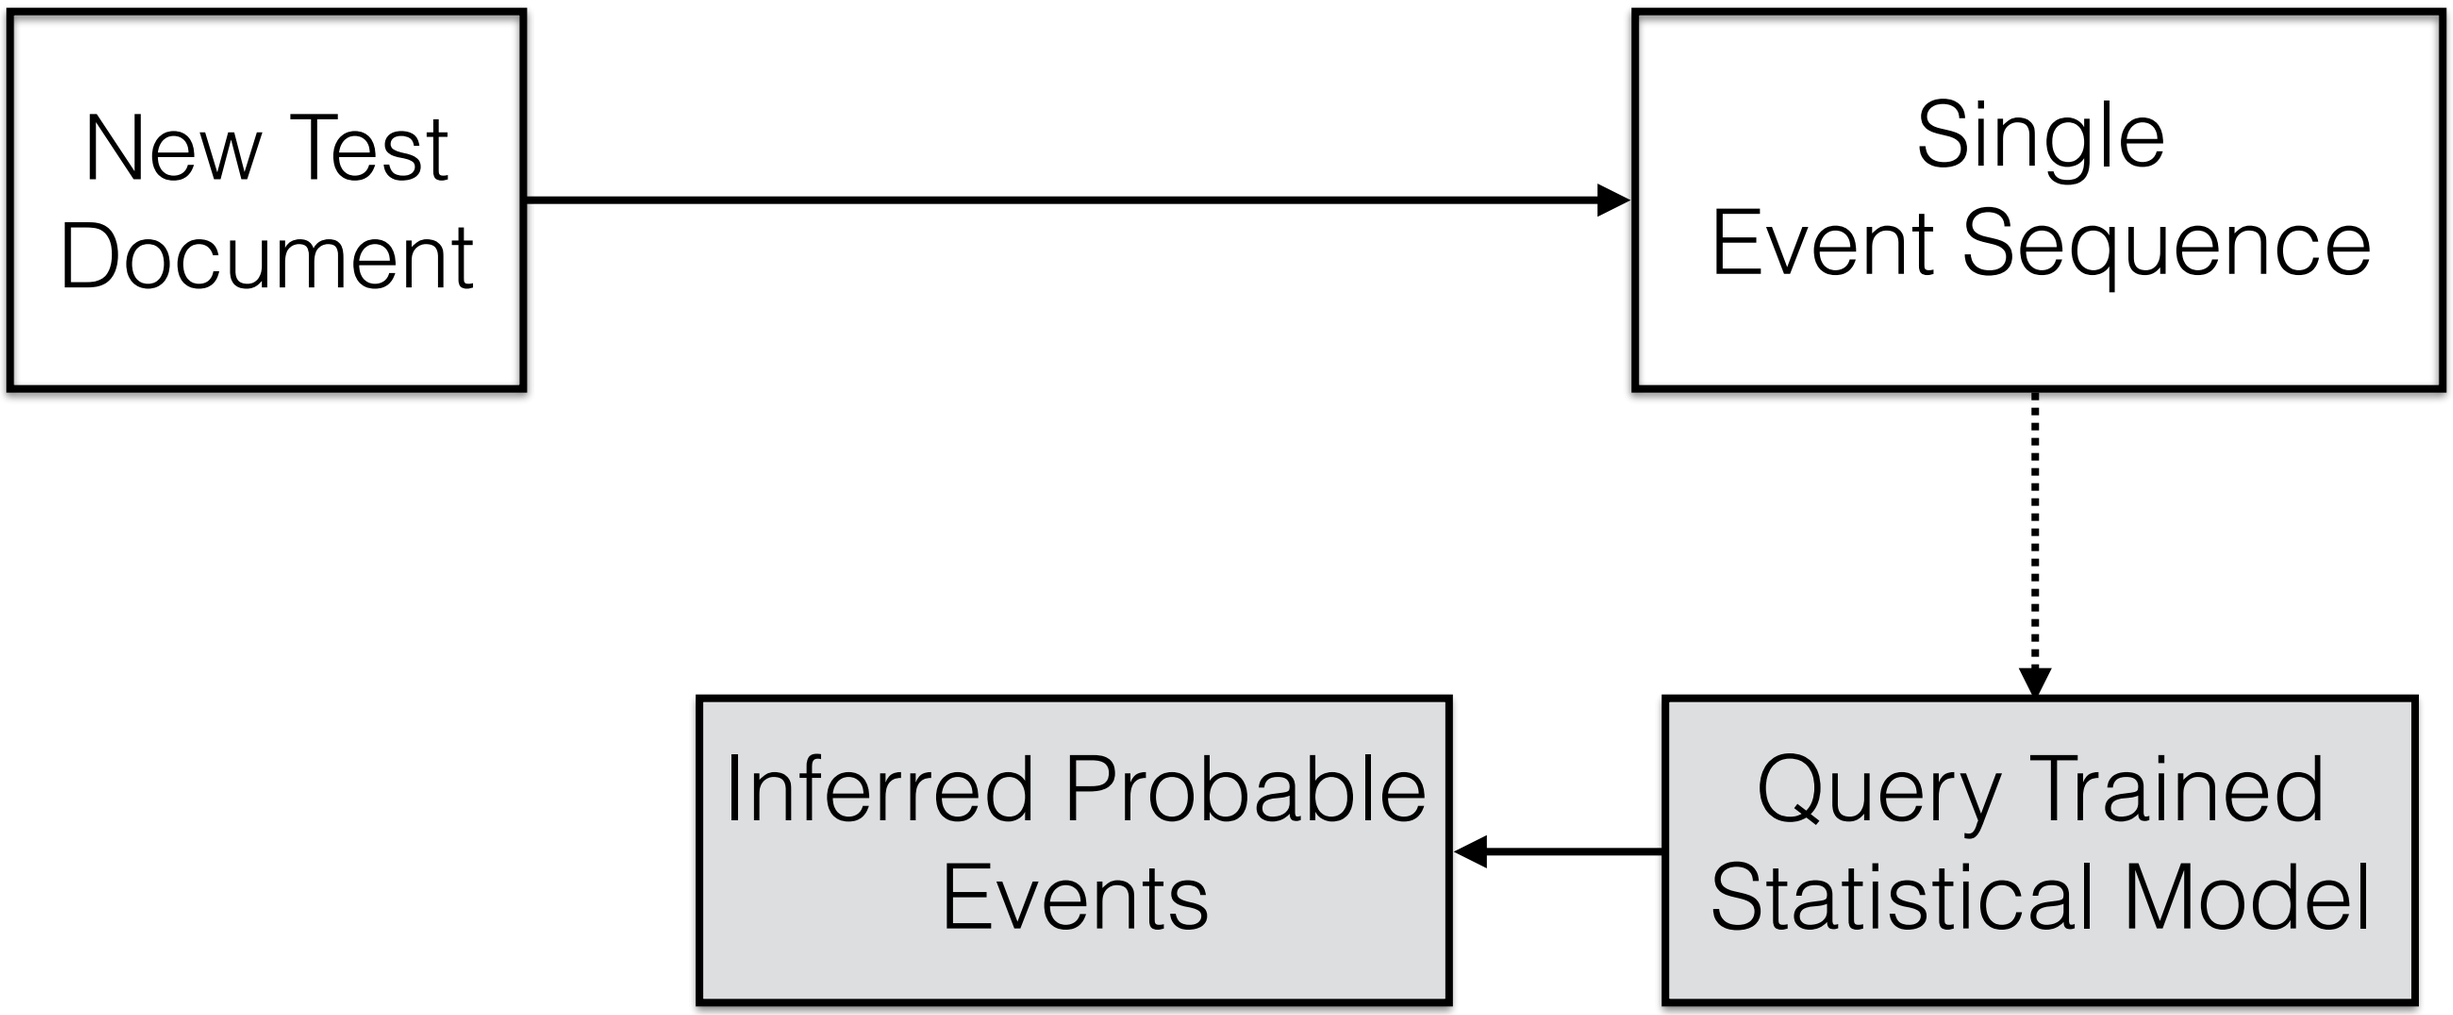

Event Sequence Inference

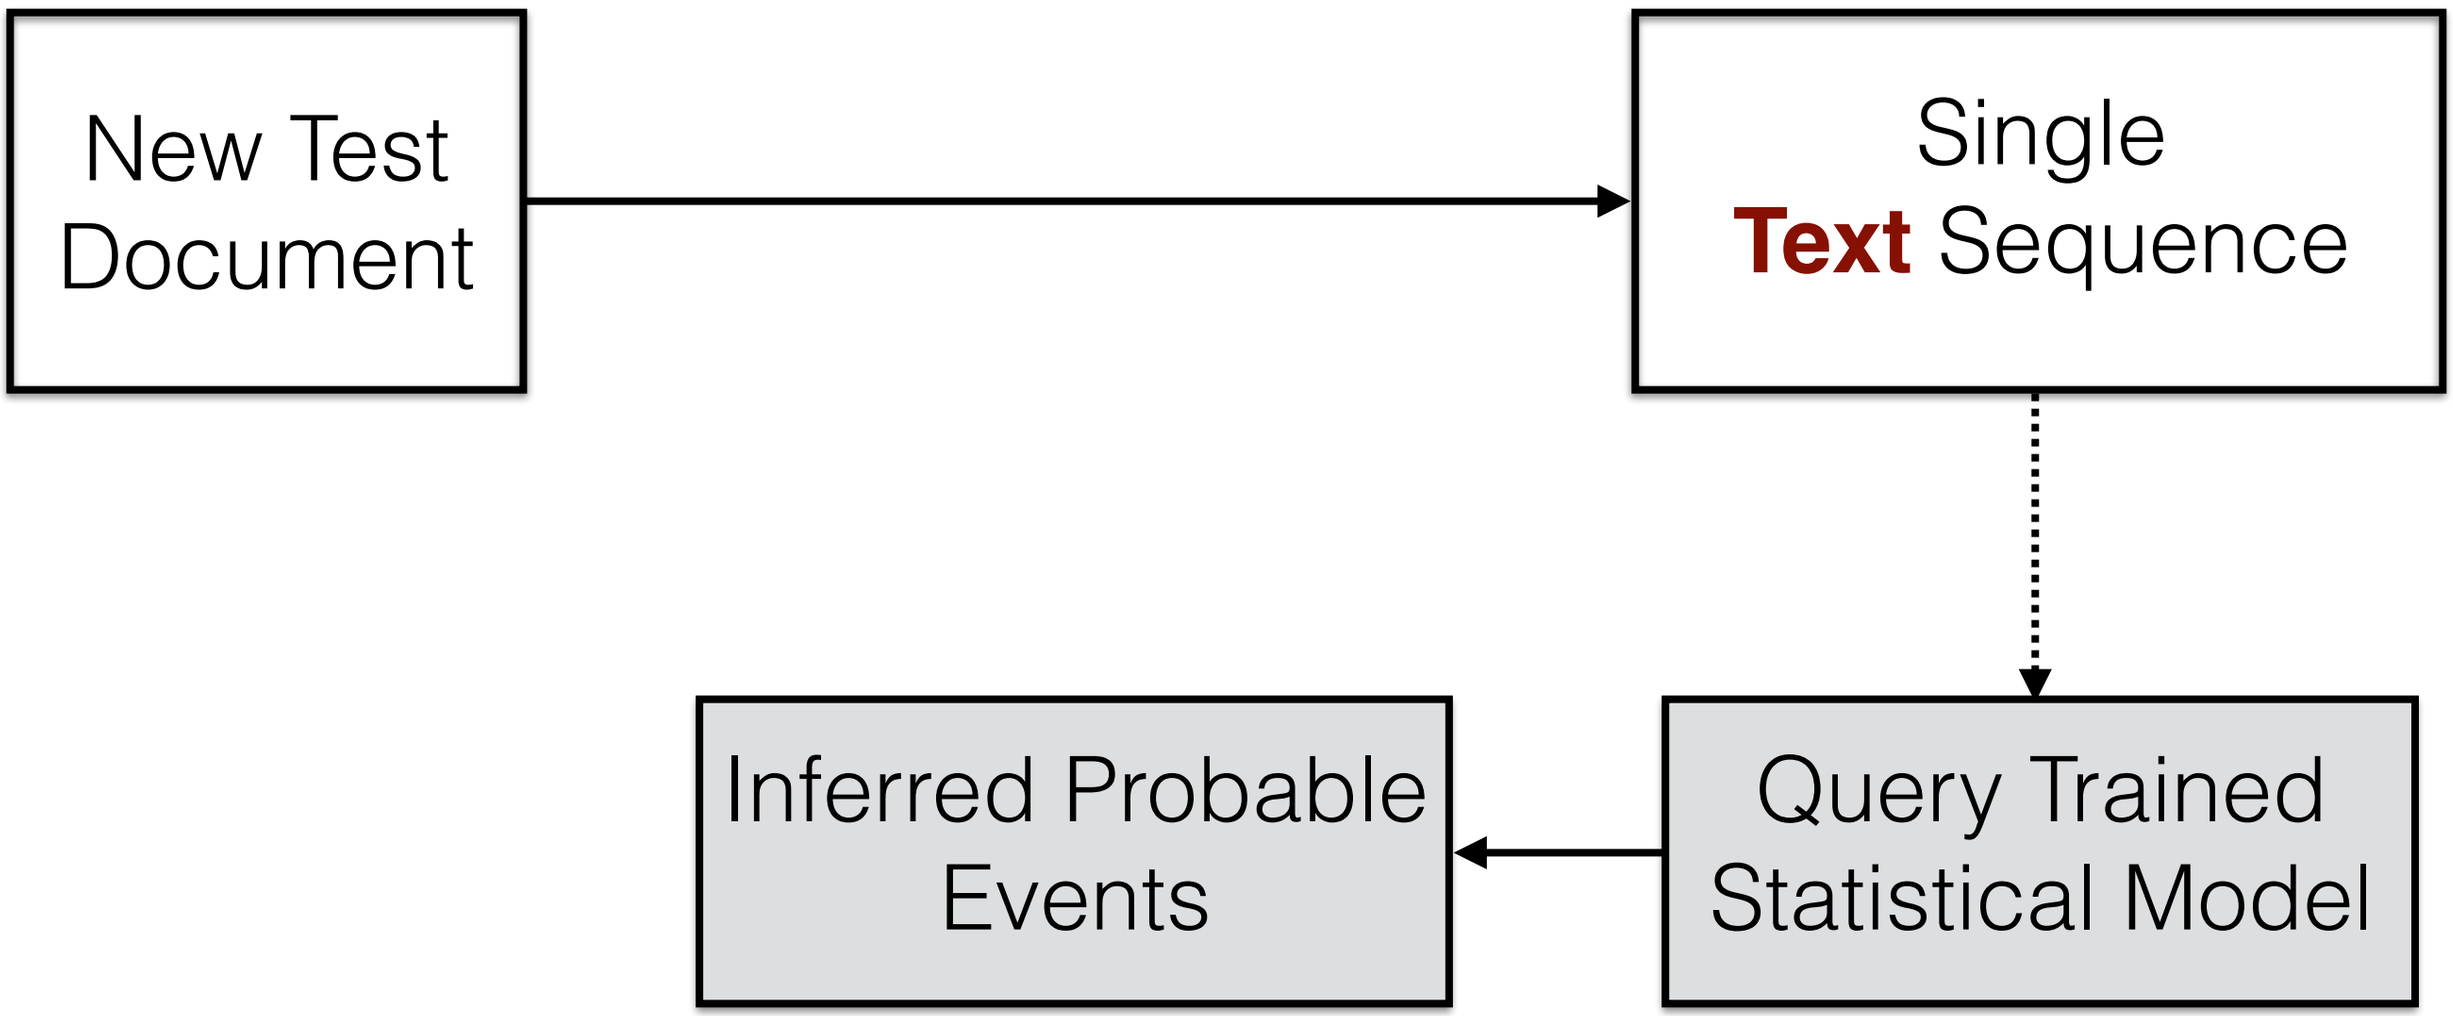

Event Sequence Inference

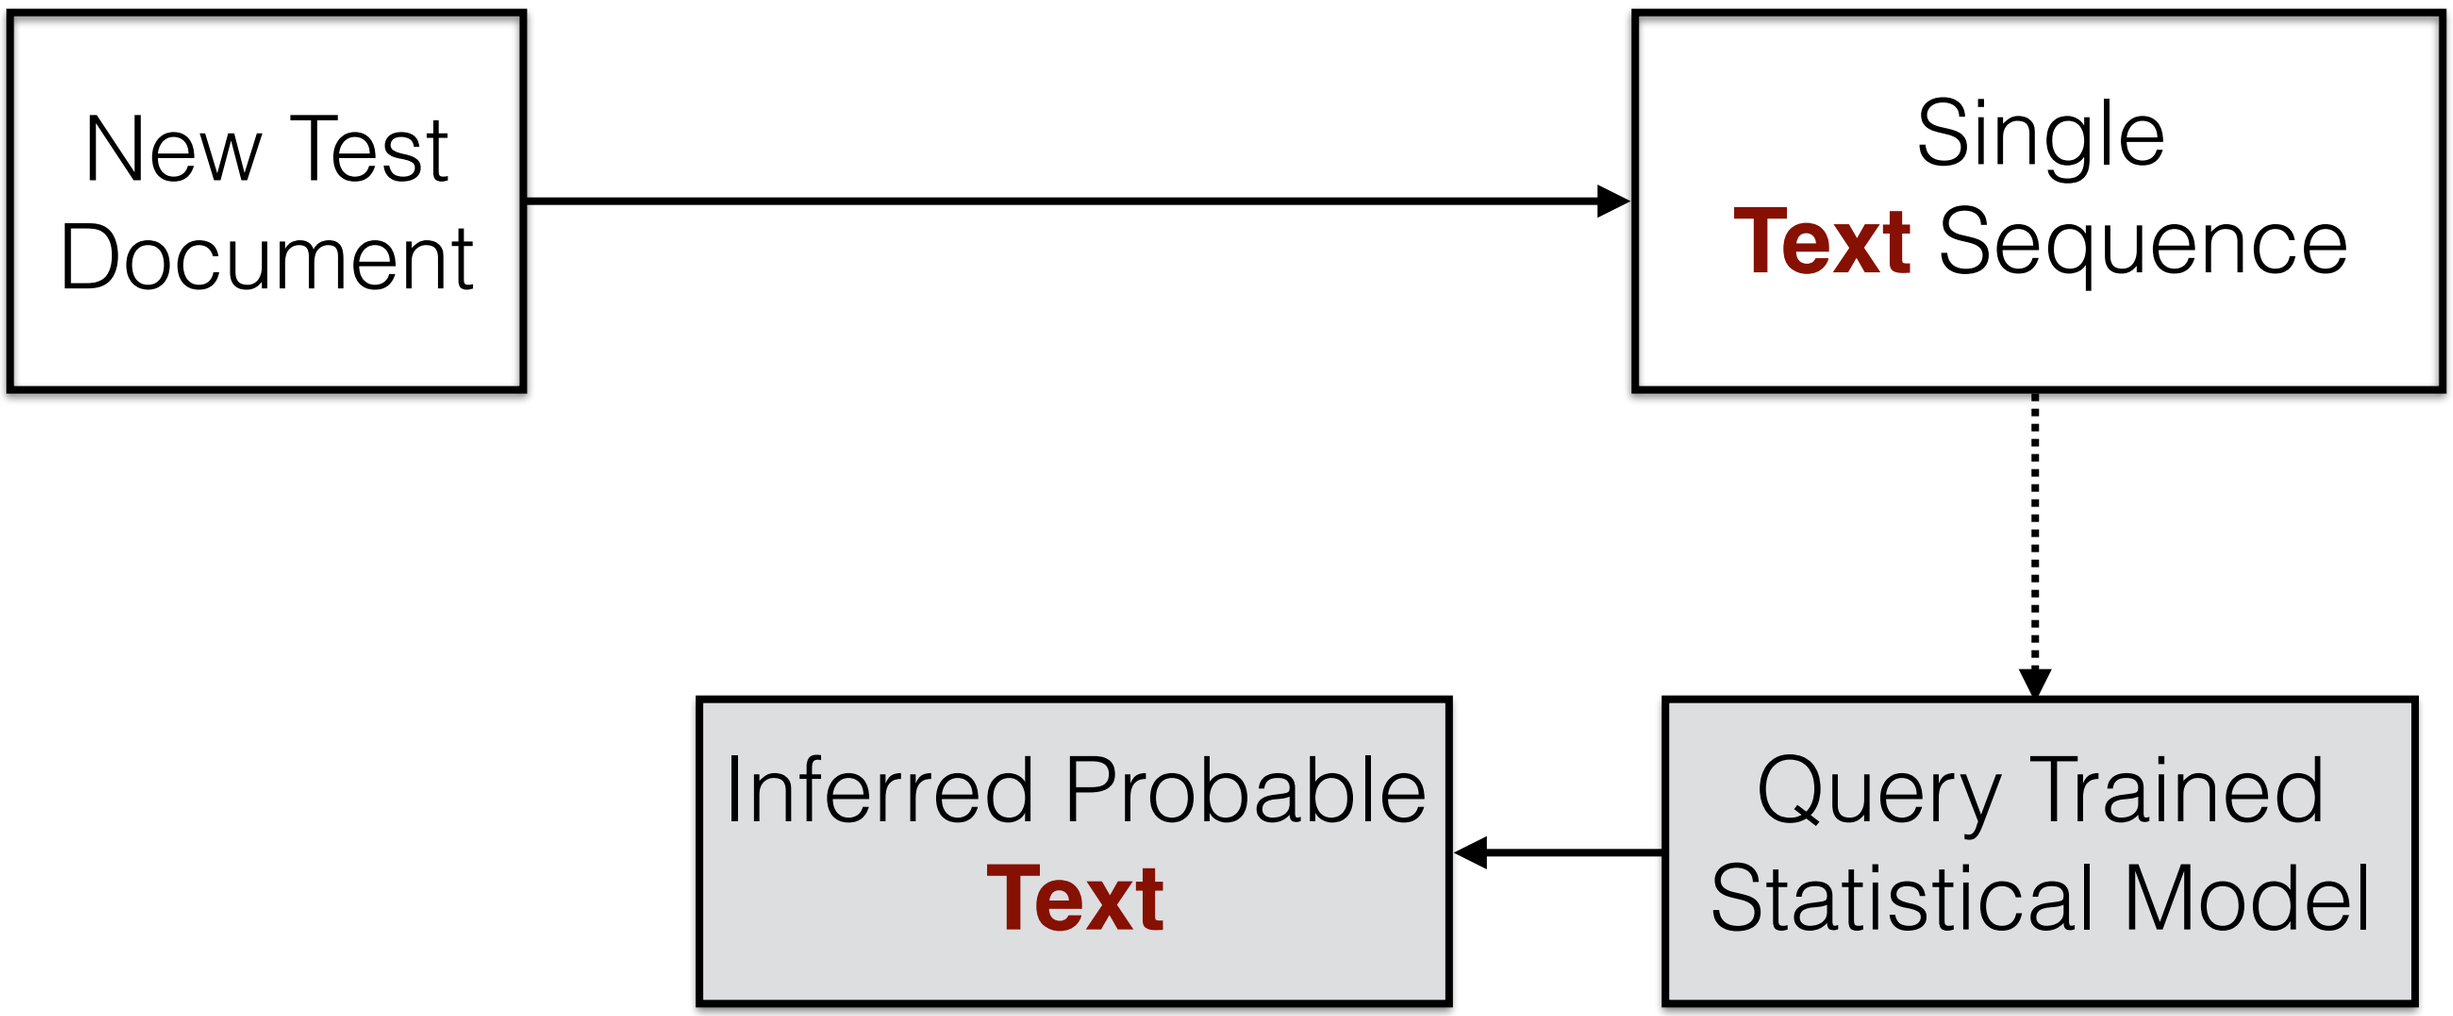

Event Sequence Inference

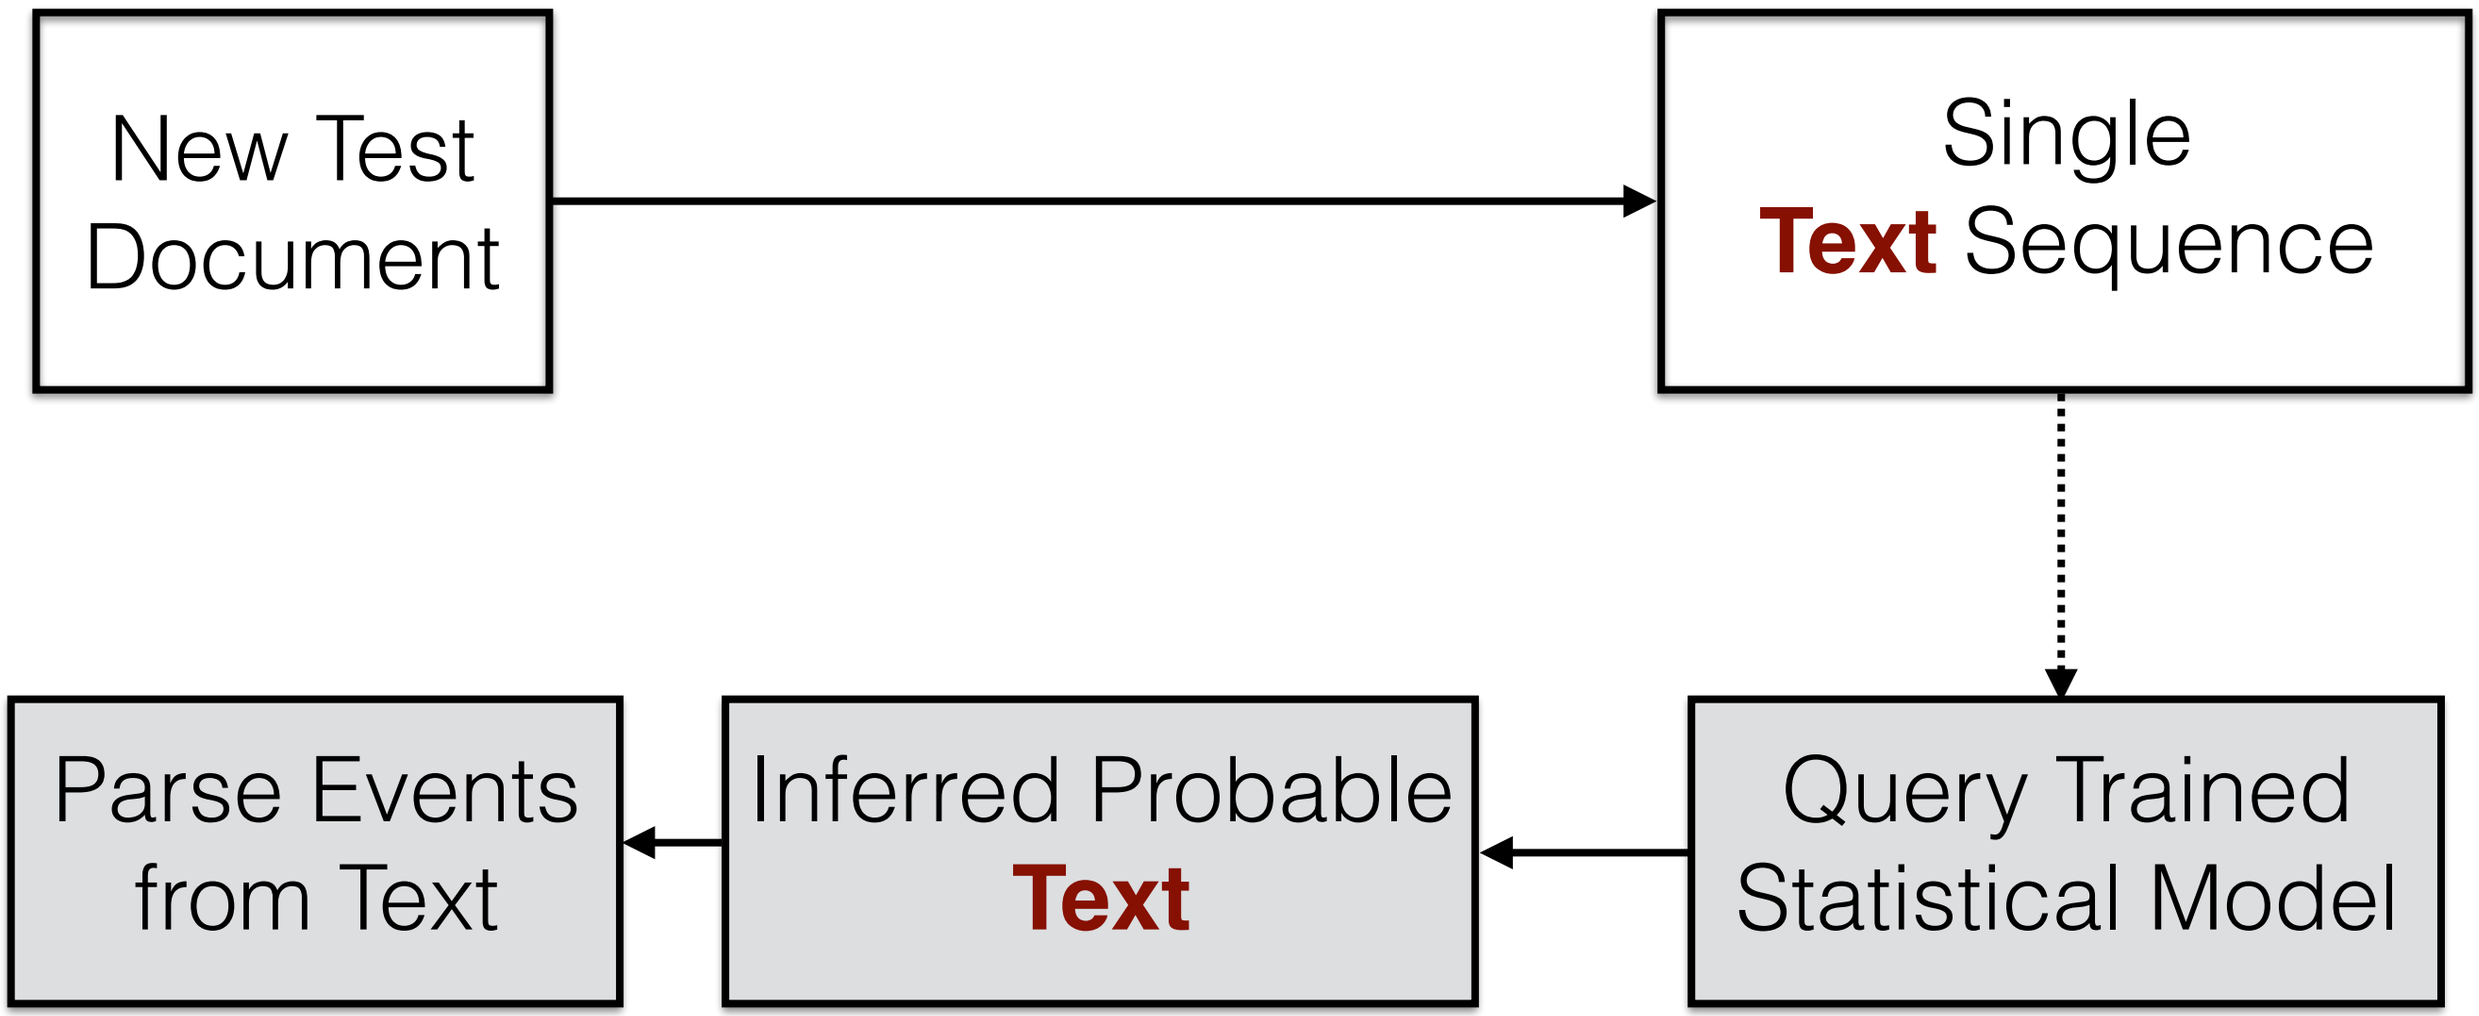

Event Sequence Inference

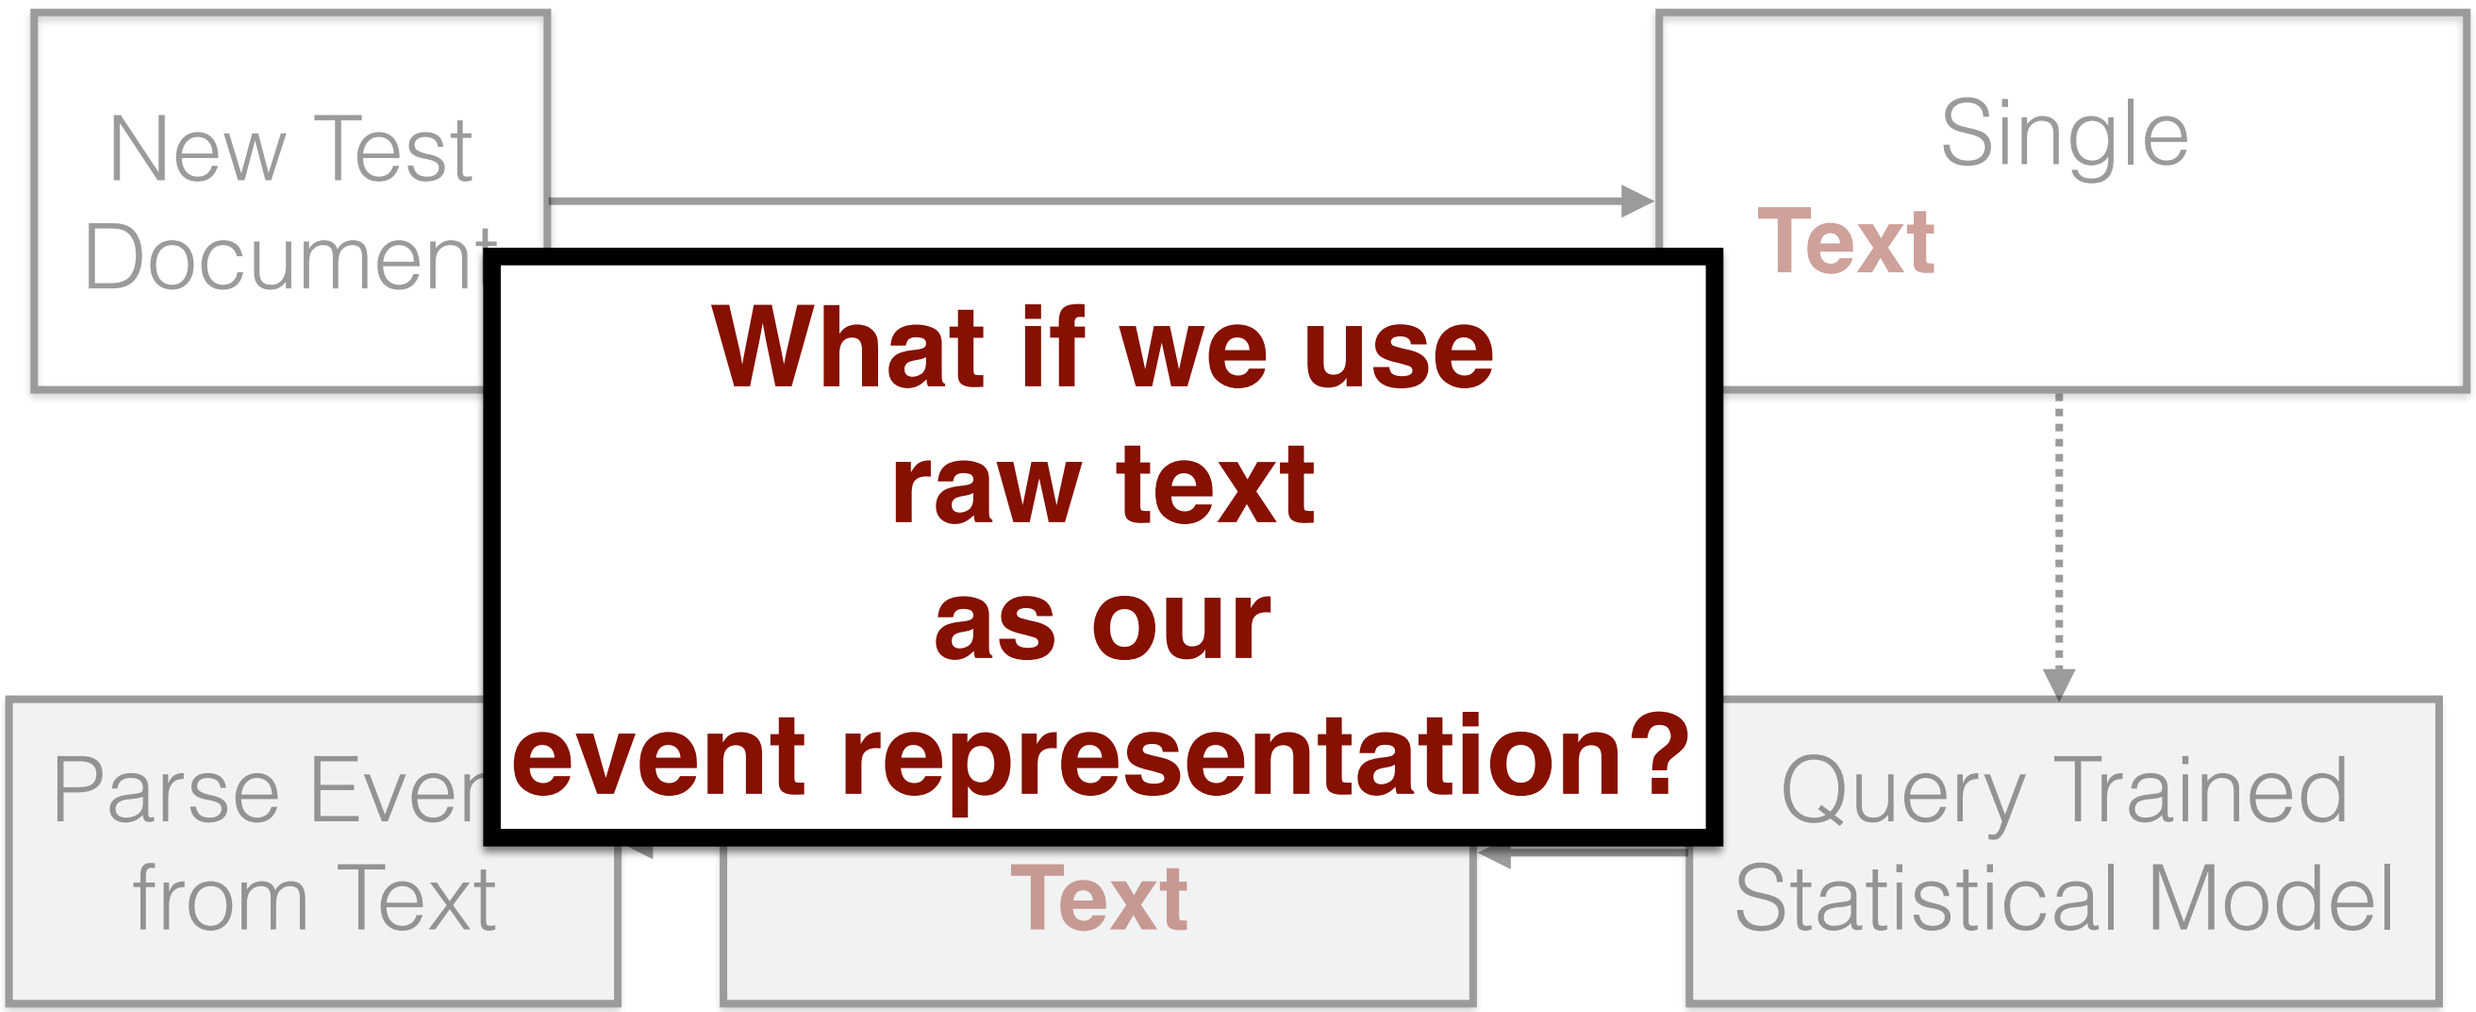

Outline

- Background & Methods
- Event Sequence Learning
- Sentence-Level Language Models

Outline

- Background & Methods
- Event Sequence Learning
- Sentence-Level Language Models

Sentence-Level Language Models

- [Kiros et al. NIPS 2015]: “Skip-Thought Vectors”
  - Encode whole sentences into low-dimensional vectors…
  - …trained to decode previous/next sentences.

Sequence-Level Language Models

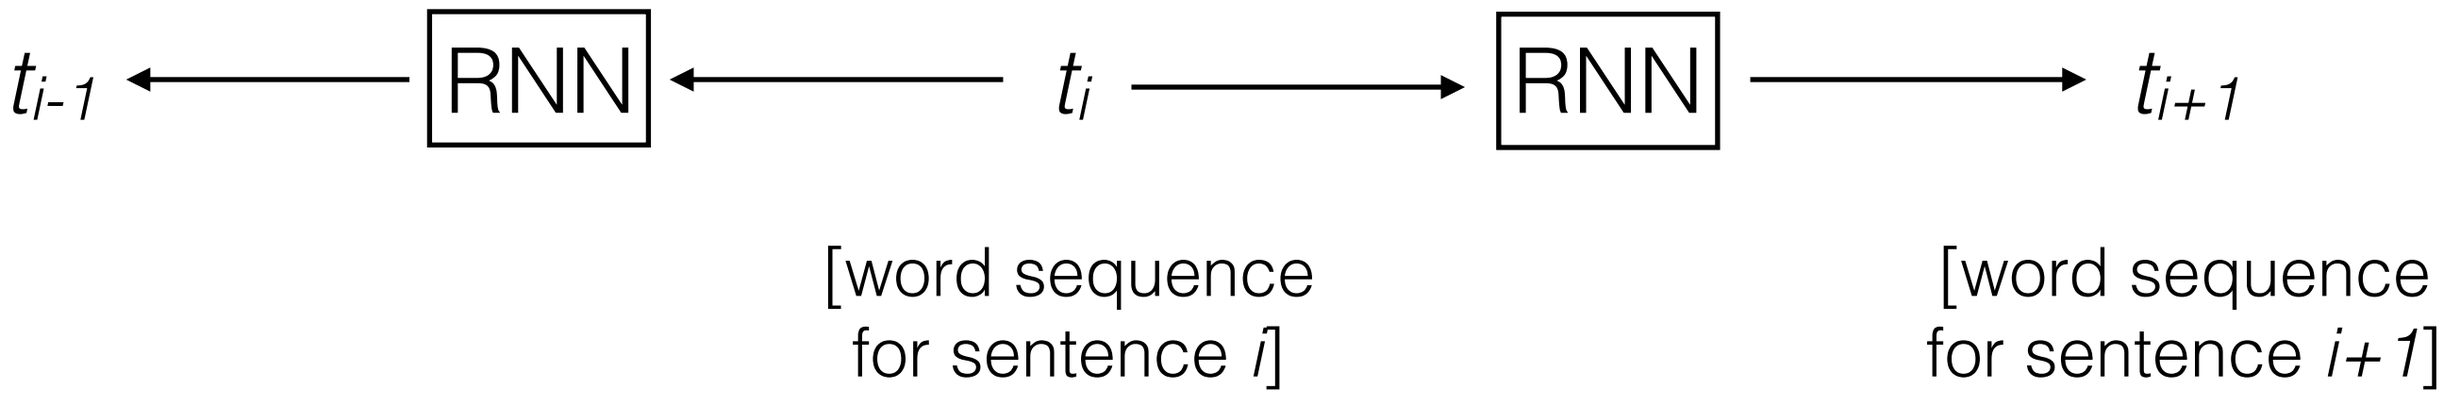

Sequence-Level Language Models

- [Kiros et al. 2015] use sentence-embeddings for other tasks.

- We use them directly for inferring text.

- **Central Question:** How well can sentence-level language models infer events?

Outline

- Background & Methods
- Event Sequence Learning
- Sentence-Level Language Models

Outline

- Background & Methods
- Experiments
- Conclusions

Outline

- Background & Methods
- Experiments
  - Task Setup
  - Results

Systems

- Two Tasks:
  - Inferring Events from Events
  - Inferring Text from Text

Systems

- Two Tasks:
  - Inferring Events from Events
    …and optionally expanding into text.
  - Inferring Text from Text
    …and optionally parsing into events.

Systems

- Two Tasks:
  - Inferring Events from Events
    …and optionally expanding into text.
  - Inferring Text from Text
    …and optionally parsing into events.
- How do these tasks relate to each other?

Event Systems

Predict an event from a sequence of events.

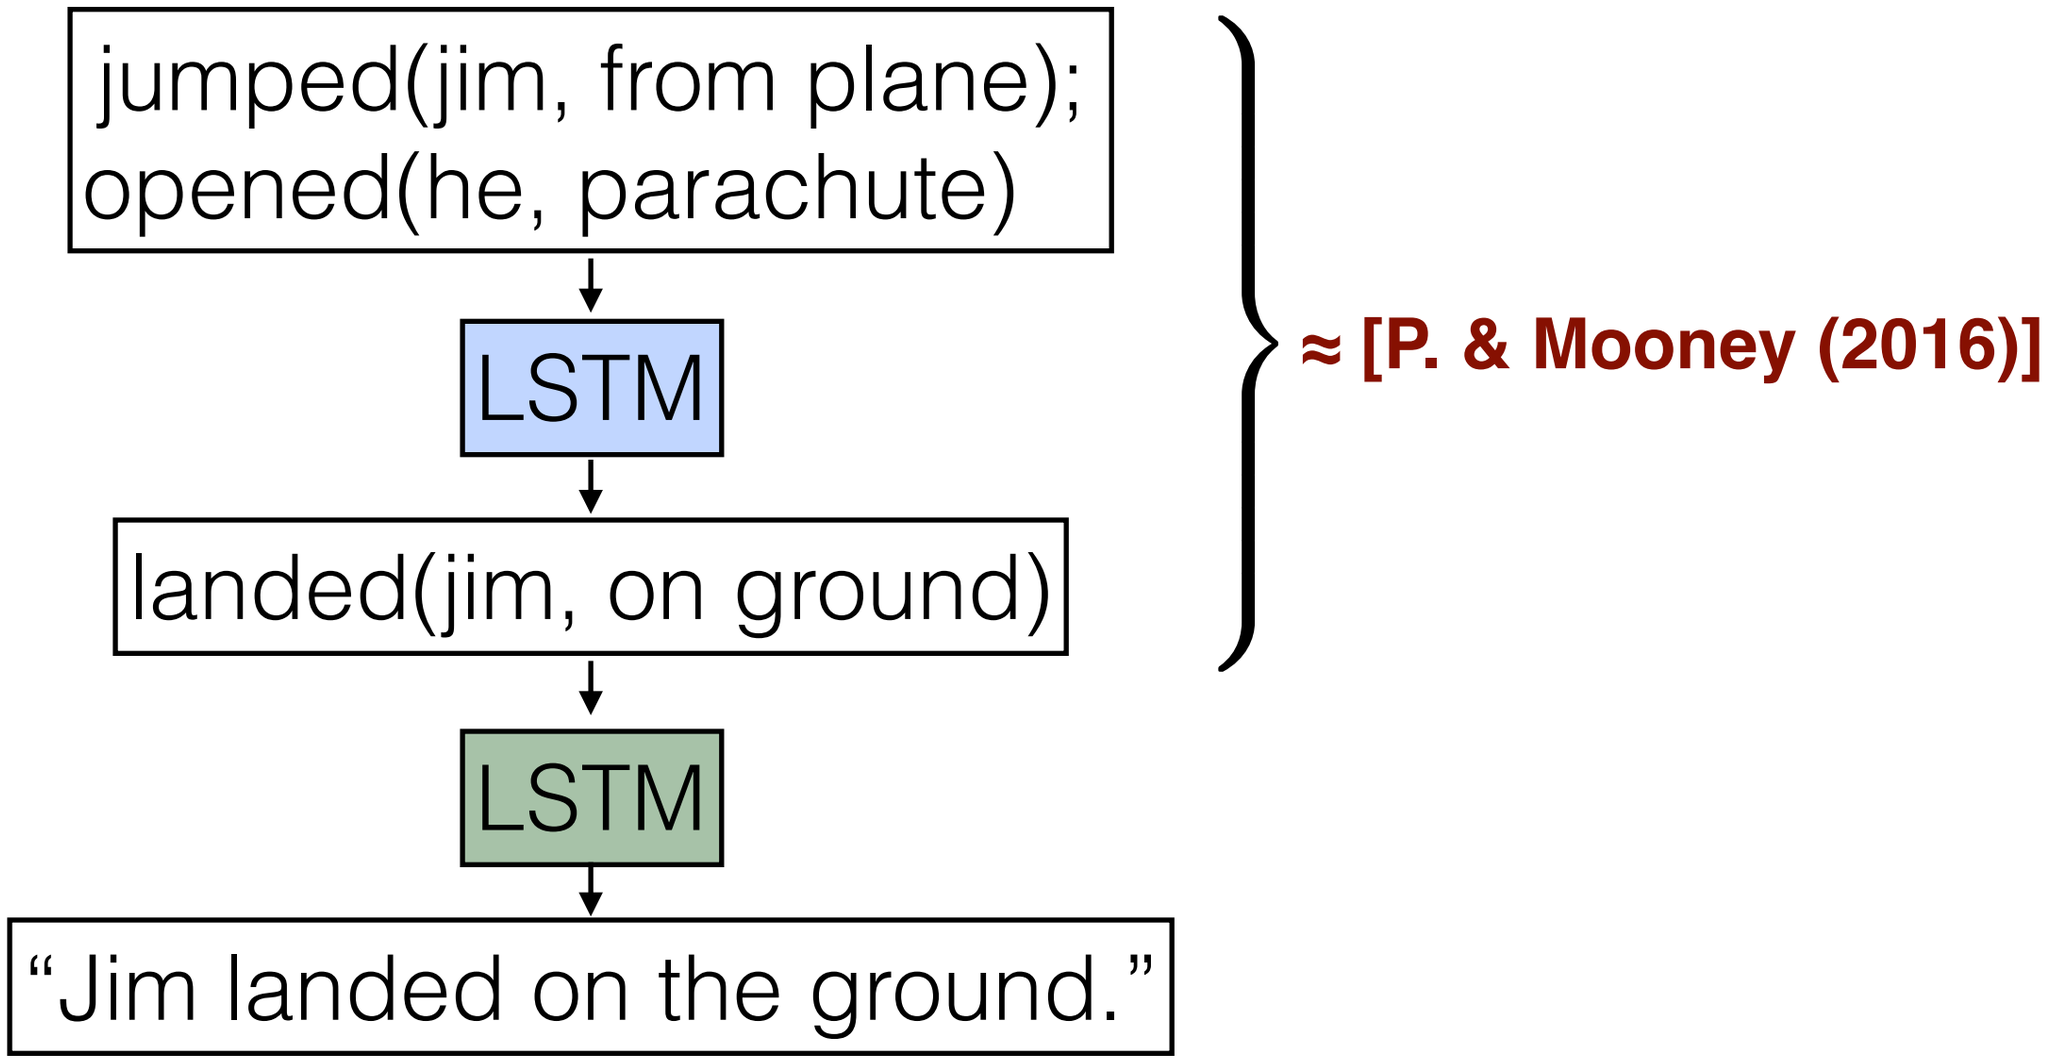

Text Systems

Predict text from text.

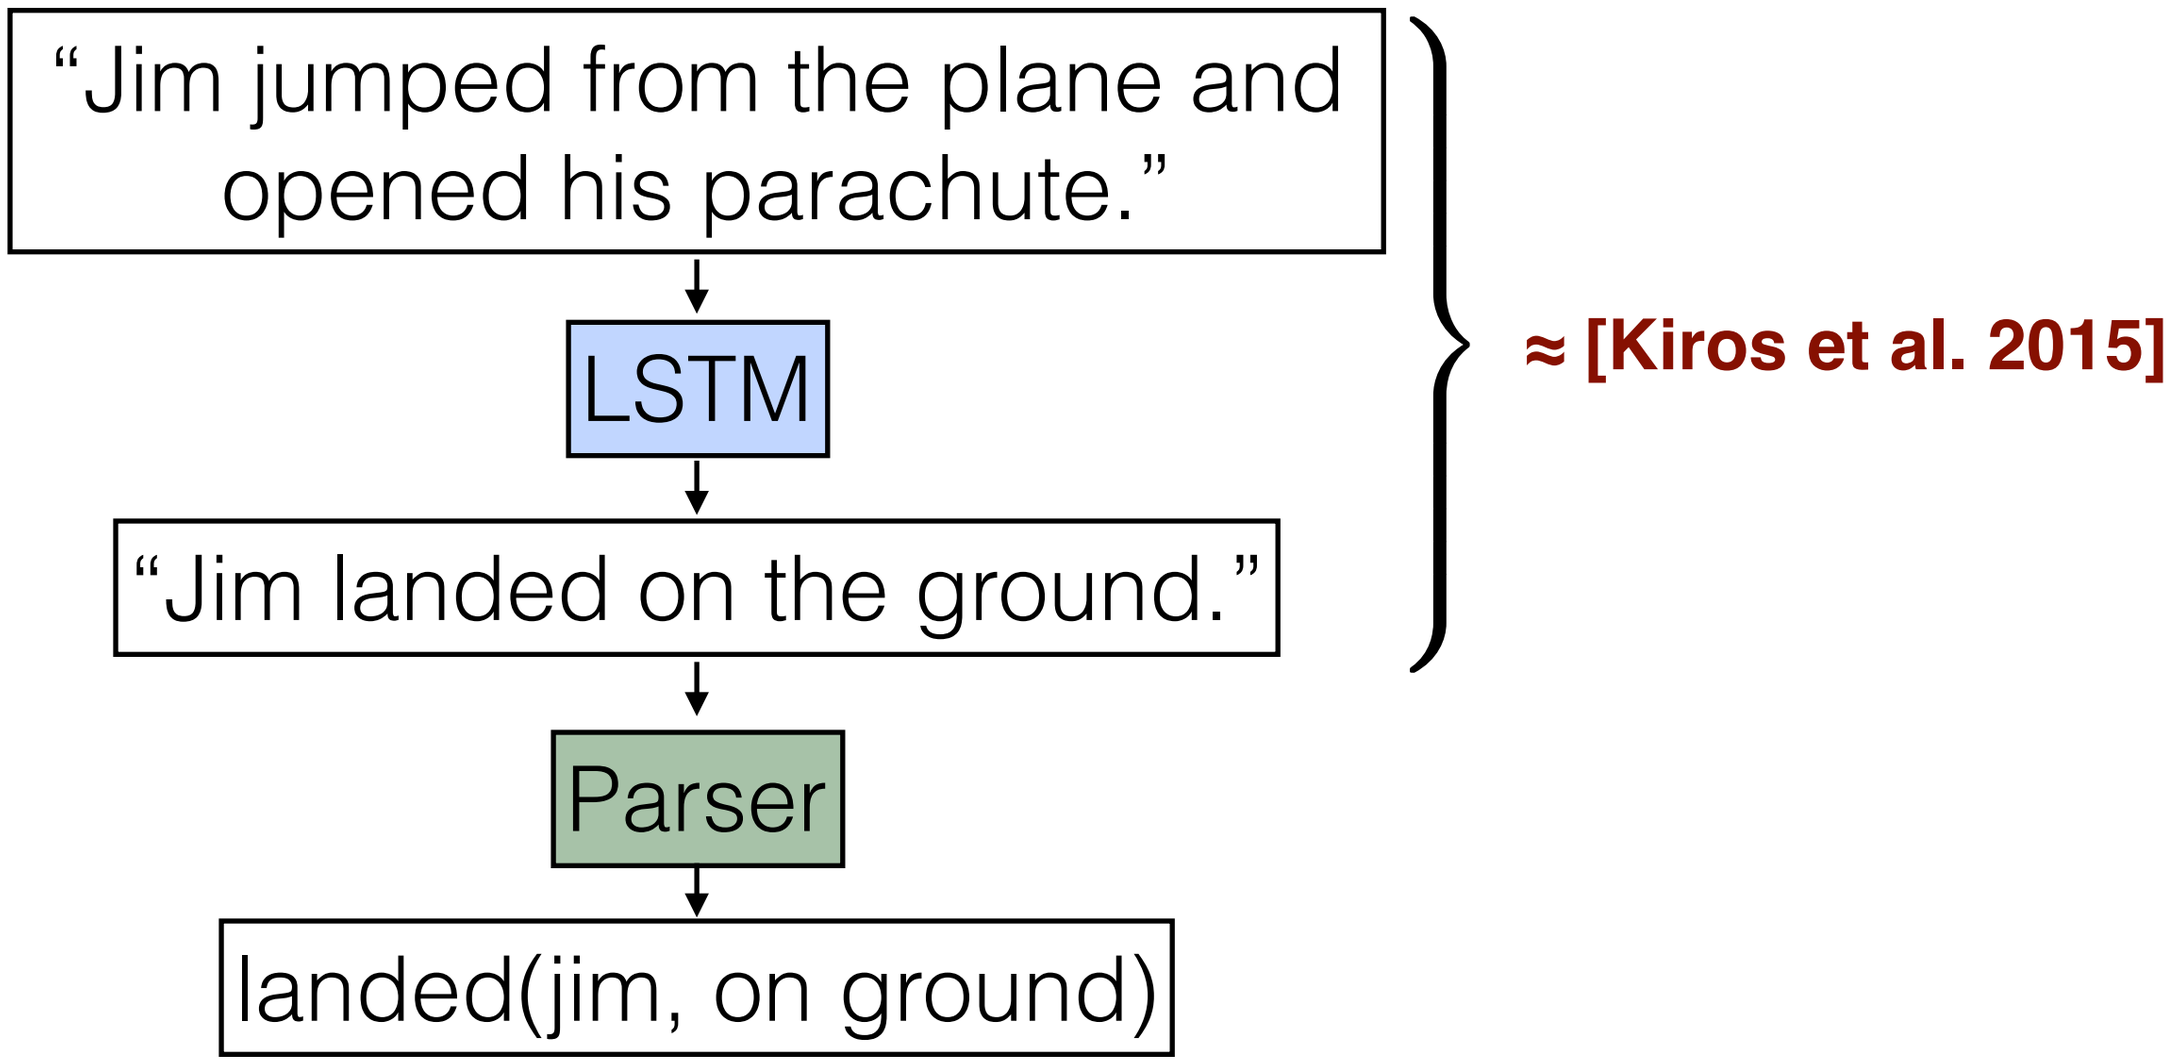

Outline

- Background & Methods
- Experiments
  - Task Setup
  - Results

Outline

- Background & Methods
- Experiments
  - Task Setup
  - Results

Experimental Setup

- Train + Test on English Wikipedia.
- LSTM encoder-decoders trained with batch SGD with momentum.
- Parse events with Stanford CoreNLP.
- Events are verbs with head noun arguments.
- Evaluate on Event Prediction & Text Prediction.

Predicting Events: Evaluation

- **Narrative Cloze** [Chambers & Jurafsky 2008]**:** Hold out an event, judge a system on inferring it.

- **Accuracy:** “For what percentage of the documents is the top inference the gold standard answer?”

- **Partial credit:** “What is the average percentage of the components of argmax inferences that are the same as in the gold standard?”

Predicting Events: Systems

- **Most Common:** Always guess the most common event.
- **e1 -> e2:** events to events.
- **t1 -> t2 -> e2:** text to text to events.

Results: Predicting Events

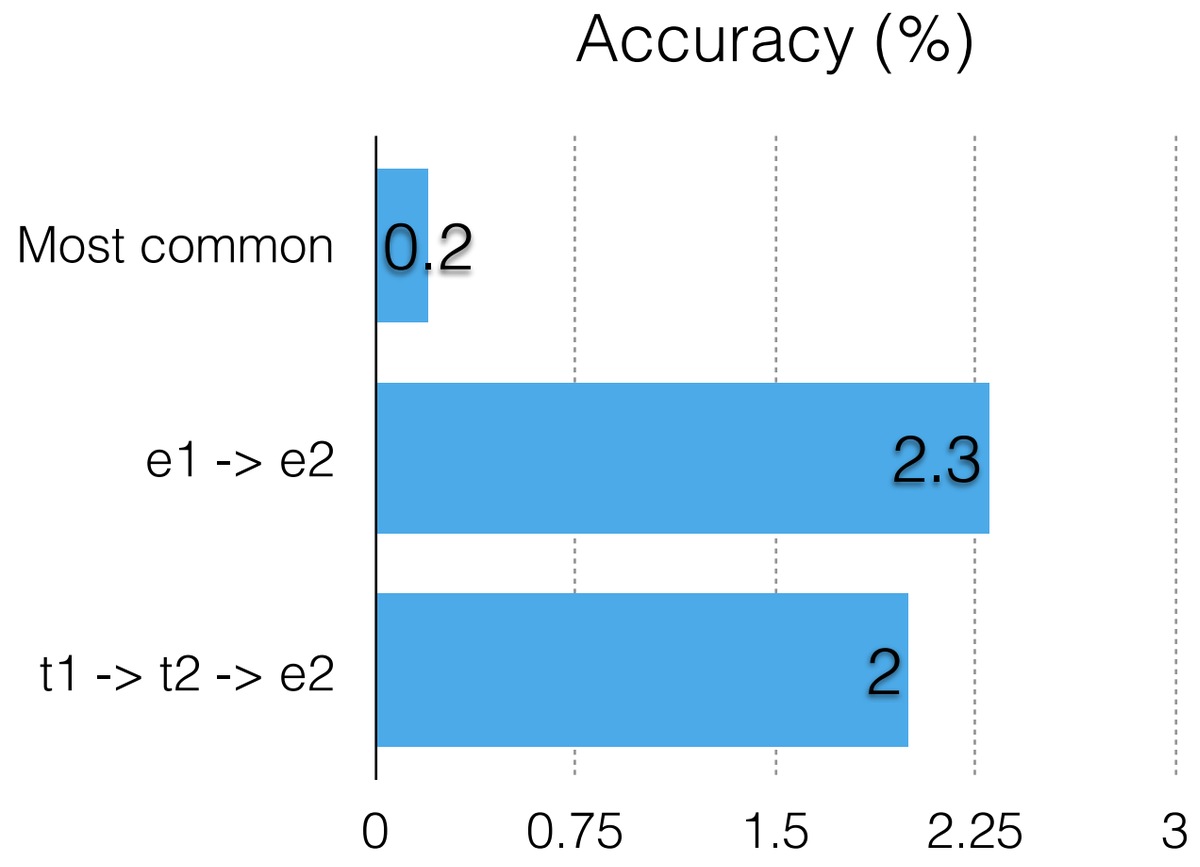

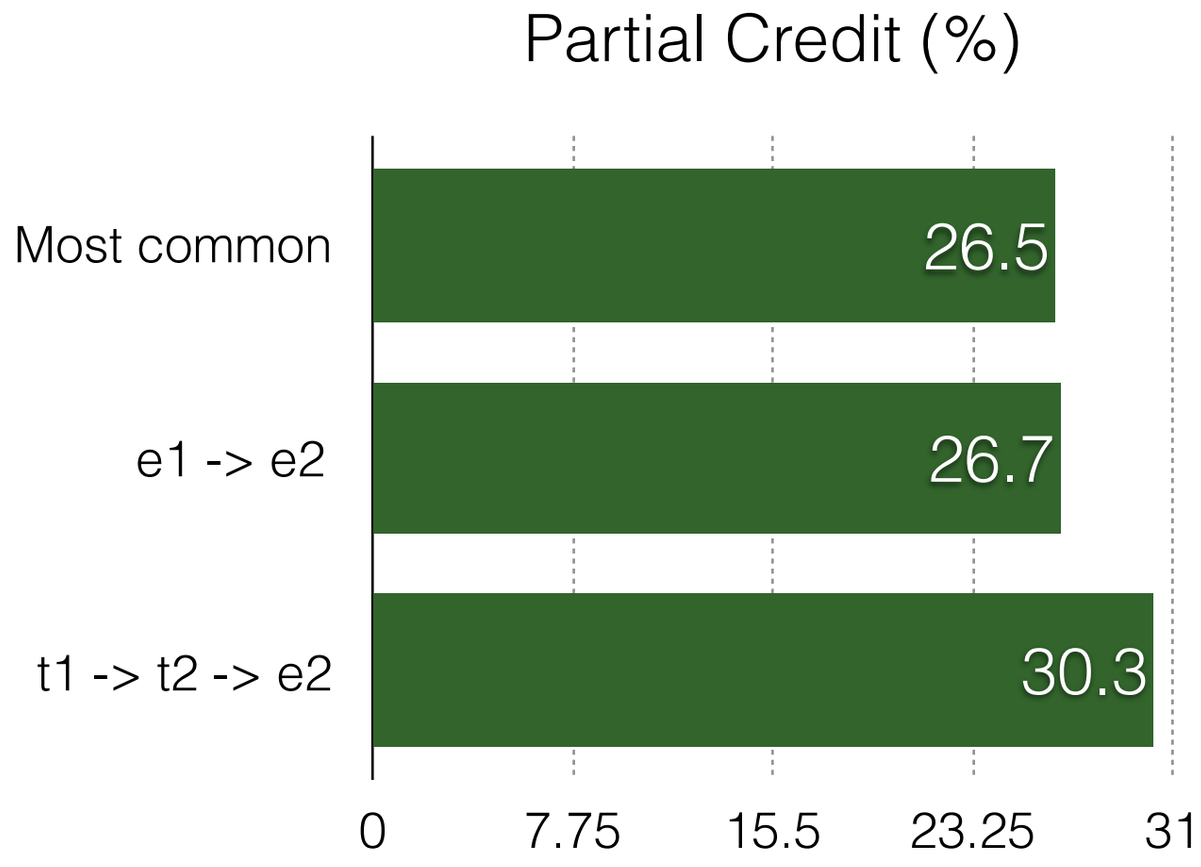

Predicting Text: Evaluation

- **BLEU:** Geometric mean of modified ngram precisions.

- Word-level analog to Narrative Cloze.

Predicting Text: Systems

- **t1 -> t1:** Copy/paste a sentence as its predicted successor.

- **e1 -> e2 -> t2:** events to events to text.

- **t1 -> t2:** text to text.

Results: Predicting Text

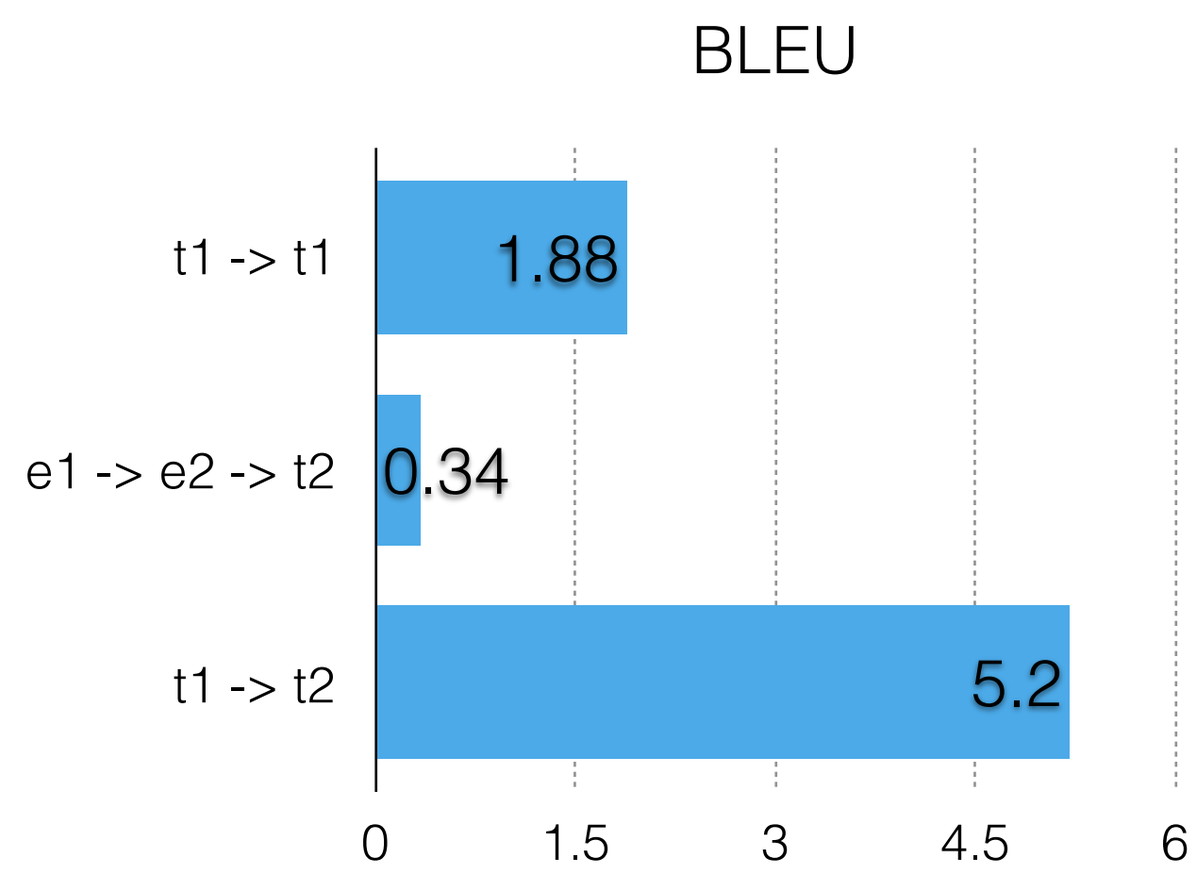

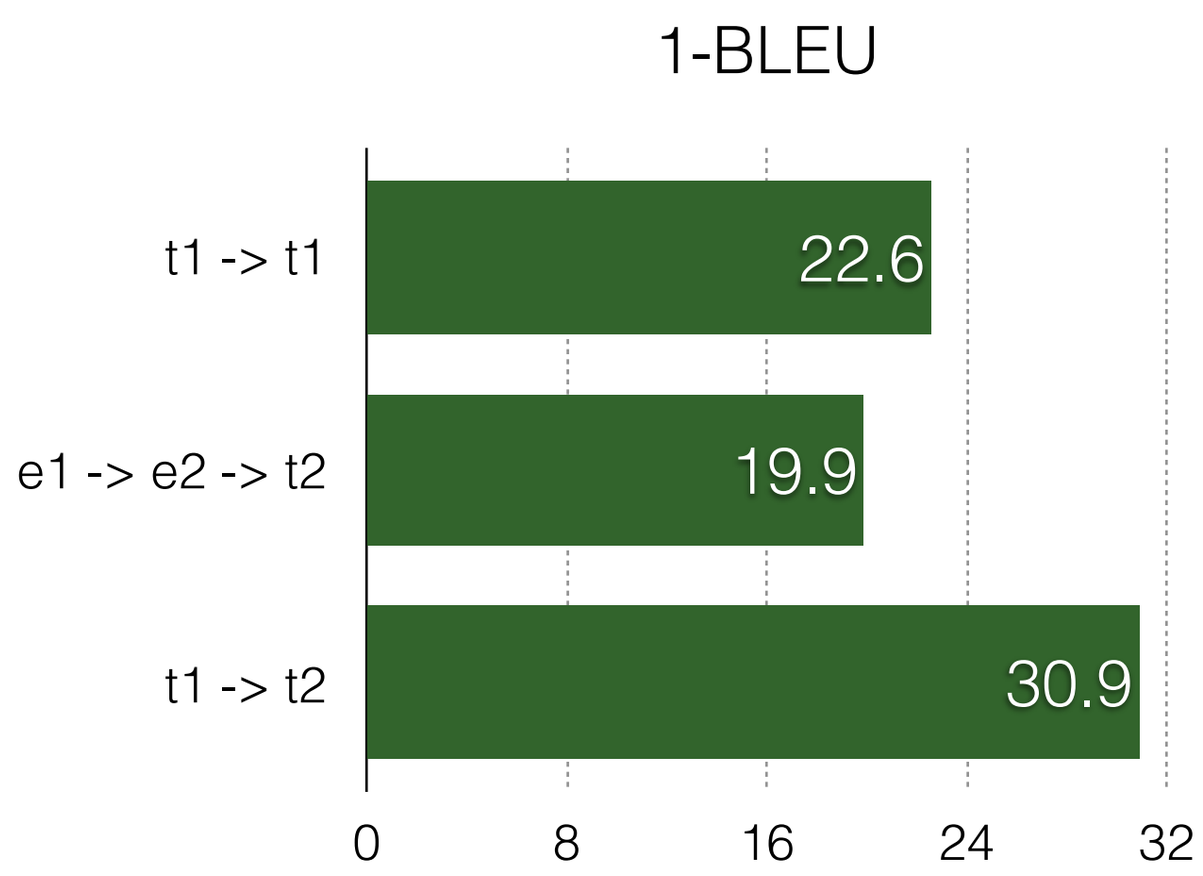

Takeaways

- In LSTM encoder-decoder event prediction…
  - Raw text models predict events about as well as event models.
  - Raw text models predict tokens better than event models.

Example Inferences

- **Input:** “White died two days after Curly Bill shot him.”

- **Gold:** “Before dying, White testified that he thought the pistol had accidentally discharged and that he did not believe that Curly Bill shot him on purpose.”

- **Inferred:** “He was buried at <UNK> Cemetery.”

Example Inferences

- **Input:** “As of October 1 , 2008 , <UNK> changed its company name to Panasonic Corporation.”

- **Gold:** “<UNK> products that were branded ‘National’ in Japan are currently marketed under the ‘Panasonic’ brand.”

- **Inferred:** “The company’s name is now <UNK>.”

Conclusions

- For inferring events in text, text is about as good a representation as events (and doesn't require a parser!).

- Relation of sentence-level LM inferences to other NLP tasks is an exciting open question.

Thanks!