From Unstructured Text to Causal Knowledge Graphs: A Transformer-Based Approach

November 2021

Scott Friedman, Ian Magnusson, Vasanth Sarathy, & Sonja Schmer-Galunder

friedman@sift.net

Today's Talk

SIFT

Cognitive Systems must reason about causality across domains.

- In this work, we focus on extracting causal knowledge from text.
- *Not* a cognitive model itself; *Not* cognitively-plausible NLP operations.
- NLP outputs support qualitative causal reasoning by cognitive systems.

Claims:

- Qualitative causal relations (and reified causal events) can capture domain-general monotonic, intentional, temporal, and functional causality.
- Transformer-based NLP can extract these causal structures from text.
- …albeit with weaker semantics than some symbolic NLP (see paper for related work!).

Results in Two Domains:

1. Scientific Claims.
2. Ethnographic Modeling.

This is just a 10-minute teaser; read the paper for details.

Causal-(ish) Language

SIFT

Scientific Journal Articles

Hedging.

Probably decrease?

Guidelines developed by experts may improve the treatment of COVID-19.

Increase or Decrease?

Levels of social support for medical staff were significantly associated with self - efficacy and sleep quality and negatively associated with the degree of anxiety and stress.

Increase or Decrease?

Thanks for the sign.

Hedging on directionality.

Most importantly, there was a significant certainty × objective ambivalence interaction, B=.04, t(169)=2.28, p=.02, 95% CI: [.01 , .07].

Causal-(ish) Language

Ethnographies & Narratives

Positive, directional, event relatedness.

It's consoling that somebody is praying for you.

“When” as temporal precedence, or as a logical conditional?

In other circumstances, pastors were consulted when women could not feel foetal movements.

Rationale, or just temporal precedence?

At 7 months I could not feel my baby move so I went to the hospital.

Possibly related by agent *intention.*

Hauwa had gone to a health facility to deliver.

Expressing Diverse Causality

Qualitative proportionalities (*q+, q-*).

- Expresses direction of causality and direction of change.
- $\alpha_{Q+/-}$ (Forbus, 1984); $M^{+/-}$ (Kuipers, 1986).

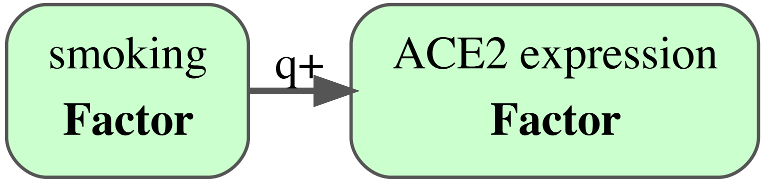

Intentions of agents, goals/subgoals (*intent+*).

- Psychological/intentional causes (Dennett, 1989).

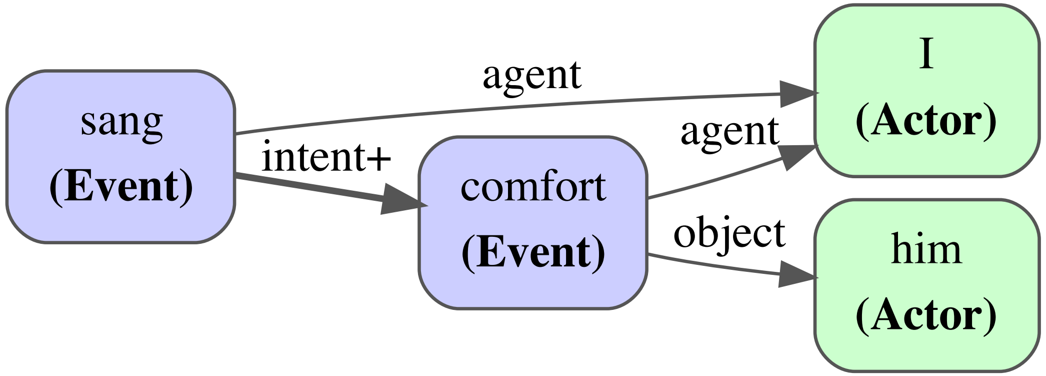

Teleology/function/design (*function+*).

- Explanation of function (Lombrozo & Carey, 2006).
- Telic affordances (Pustejovsky, 1991).

Temporal precedence (*t+*).

- Expresses temporal ordering.
- *May* indicate causality, norms, plans, “scripts.”

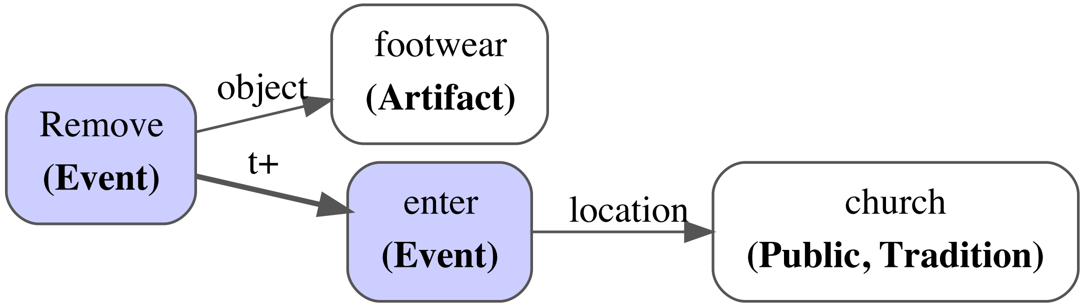

From Text to Knowledge Graphs

Web UI for Annotating Relational NLP Examples

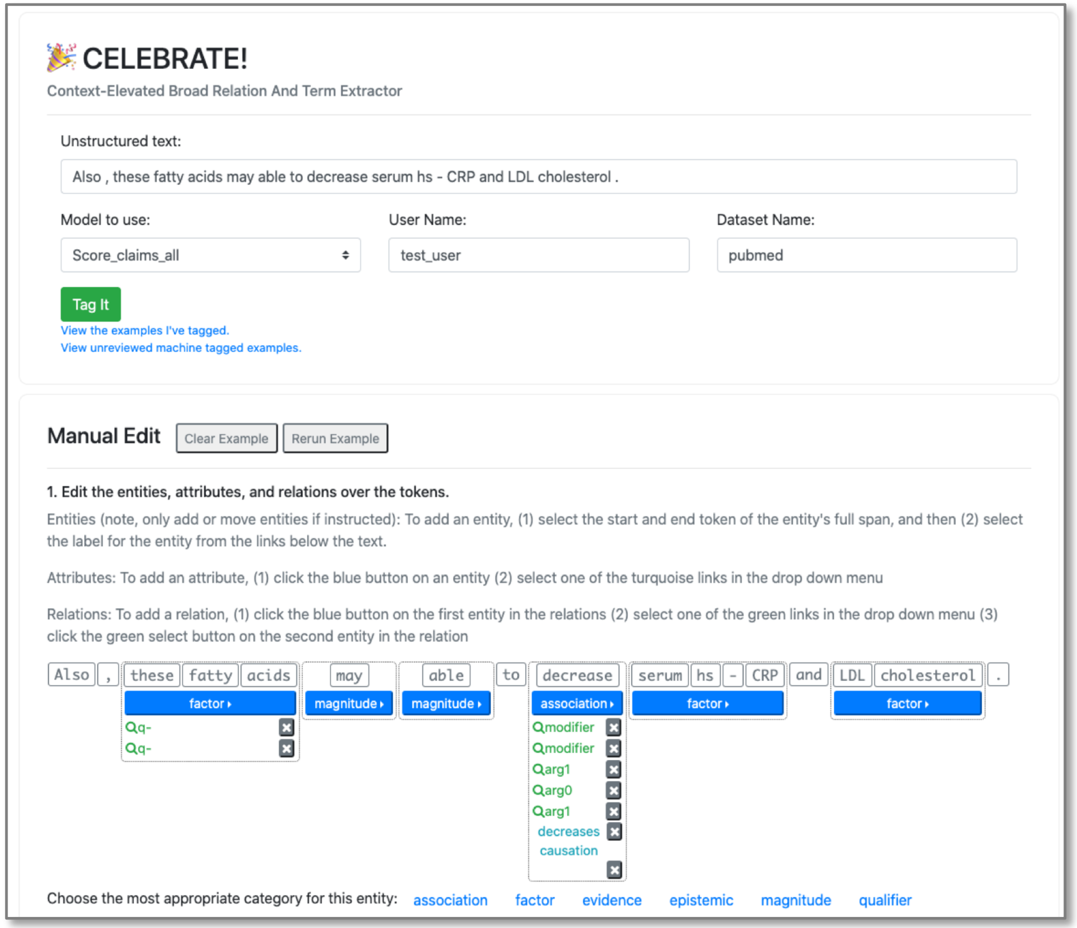

SpEAR: Transformer-based NLP Architecture for Extracting Knowledge Graphs from Text

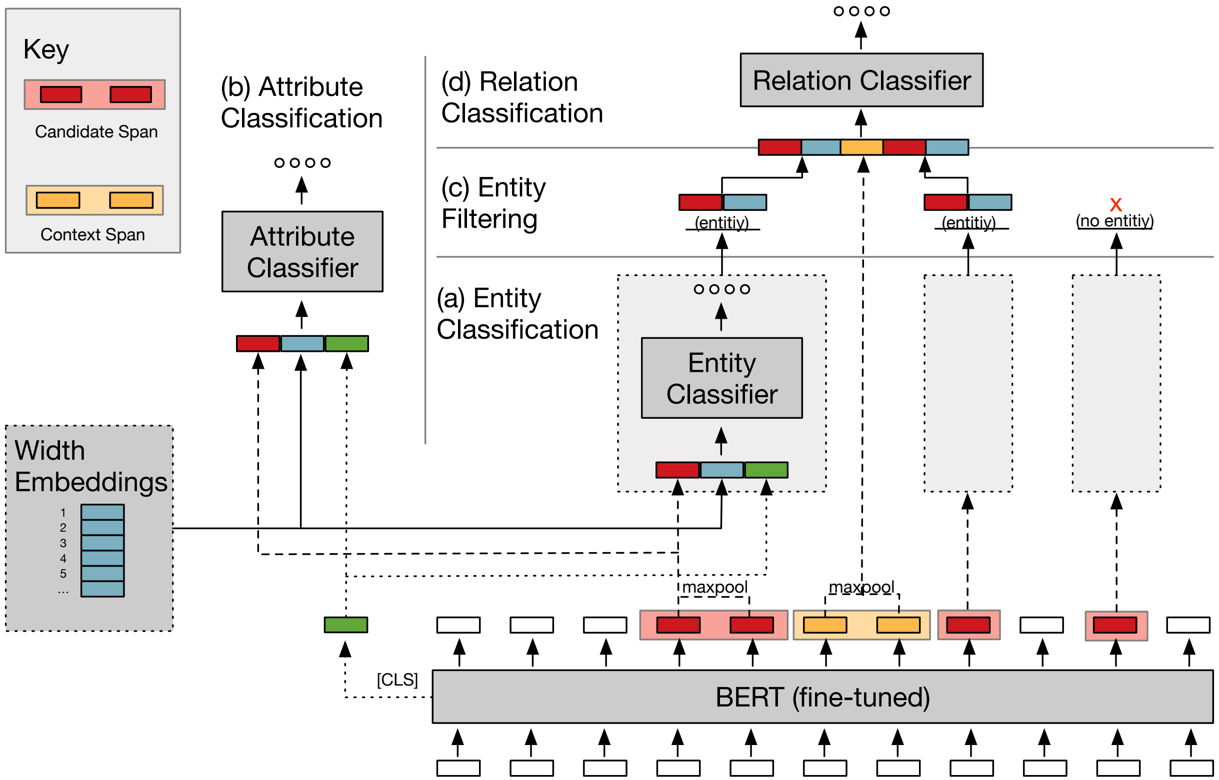

Based on SpERT:
Eberts & Ulges. (2020). Span-based joint entity and relation extraction with transformer pre-training. *ECAI 2020*.

Knowledge Graphs over Text

Input text: *"The women prayed to prevent any complications."*

Relations are labeled, directed edges.

Attributes are multi-class types expressed on nodes.

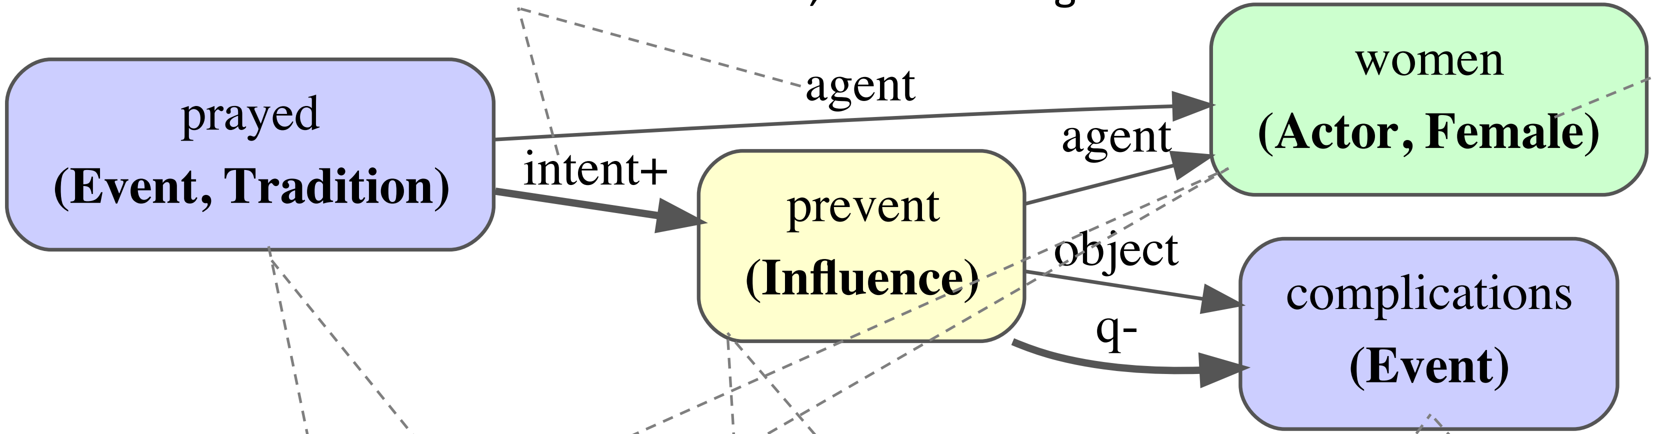

Nodes (i.e., **entities**) are contiguous, potentially-overlapping spans of tokens.

The [**women**] [**prayed**] to [**prevent**] any [**complications**] .

Present representational constraint: No nodes/entities directly inferred without token support.

WordNet Word Senses as Semantic Labels

Use LMMS (Loureiro & Jorge, 2019) to:

- Encode a vector for each token.
- Dot-product each nodes’ token(s) against WordNet sense keys.
- Compute top-*k* sense keys above a threshold (e.g., 0.5).

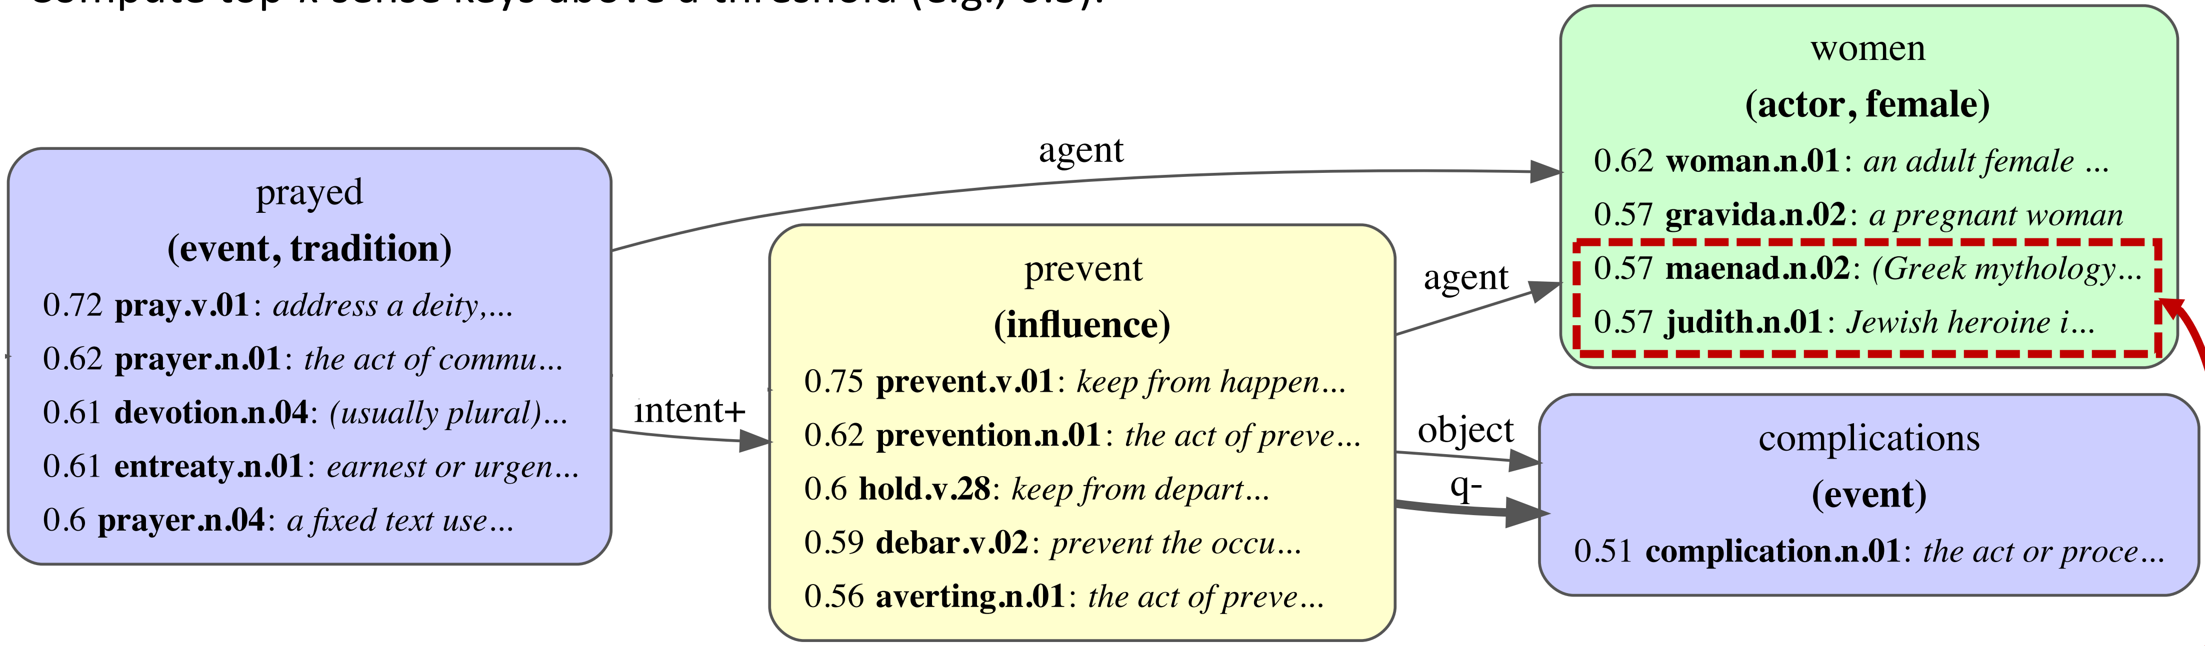

Sometimes triggers unusual word senses.

...this maps each node’s tokens into the WordNet ontology, potentially to multiple locales, weighted by confidence.

Domain 1: Scientific Claims

Currently 900 sentences from:

1. PubMed.

- Selected from Yu et al.'s *"Detecting causal language…"*

Bei Yu, Yingya Li, and Jun Wang. (2019). Detecting causal language use in science findings. *EMNLP*, p 4656–4666.

2. Social and Behavior Science (SBS) literature.

- Selected from the Center for Open Science's SCORE dataset.

Nazanin Arendt, Daniel Jacob Benjamin, Noam Benkler, Michael Bishop, Mark Burstein, Martin Bush, James Caverlee, Yiling Chen, Chae Clark, et al. (2021). Systematizing confidence in open research and evidence (SCORE).

3. CORD-19: COVID-19 Open Research Dataset.

Lucy Lu Wang, Kyle Lo, Yoganand Chandrasekhar, Russell Reas, et al. (2020). CORD-19: The COVID-19 open research dataset.

Domain 1: Results

	Dimension	P	R	F1	Support
Entities	factor	90.13	86.71	88.39	1,604
	evidence	72.73	80.00	76.19	139
	epistemic	93.33	100.00	96.55	178
	association	95.89	93.33	94.59	837
	magnitude	94.44	94.44	94.44	415
	qualifier	86.96	68.97	76.92	216
	Micro-Averaged	91.29	87.89	89.56	
Attributes	causation	88.24	93.75	90.91	204
	comparison	79.17	90.48	84.44	234
	indicates	80.00	66.67	72.73	44
	increases	75.86	95.65	84.62	262
	decreases	100.00	100.00	100.00	134
	correlation	94.74	94.74	94.74	199
	test	100.00	66.67	80.00	24
	Micro-Averaged	84.62	91.67	88.00	
Relations	arg0	82.93	76.40	79.53	865
	arg1	76.71	71.79	74.17	883
	comp_to	81.82	69.23	75.00	137
	modifier	84.78	74.29	79.19	1,080
	q+	77.78	56.00	65.12	295
	q-	60.00	85.71	70.59	138
	subtype	85.71	75.00	80.00	106
	Micro-Averaged	81.00	72.97	76.78	

Table 1: Precision, recall, F1 and support (i.e., occurrences in dataset) for each label on 10% held-out dataset using SpEAR with rectifier and filtering model.

- Highest F1 on **Entity** extraction.
- **Attributes** are close behind.
  - These depend on a correct entity extracted.
- **Relations** are a focus of near-term work.
  - These depend on *two* correct entities extracted.
  - …and representing relational context for a linear layer to infer the label and directionality.

Domain 2: Ethnographic Analysis

Currently 700+ sentences & paragraphs from:

1. Ethnographies and Anthropology journal articles.
2. Ethnographic interview transcripts.
3. Folk tales.
4. Social media.

Propagate Valence via Qualitative Causality

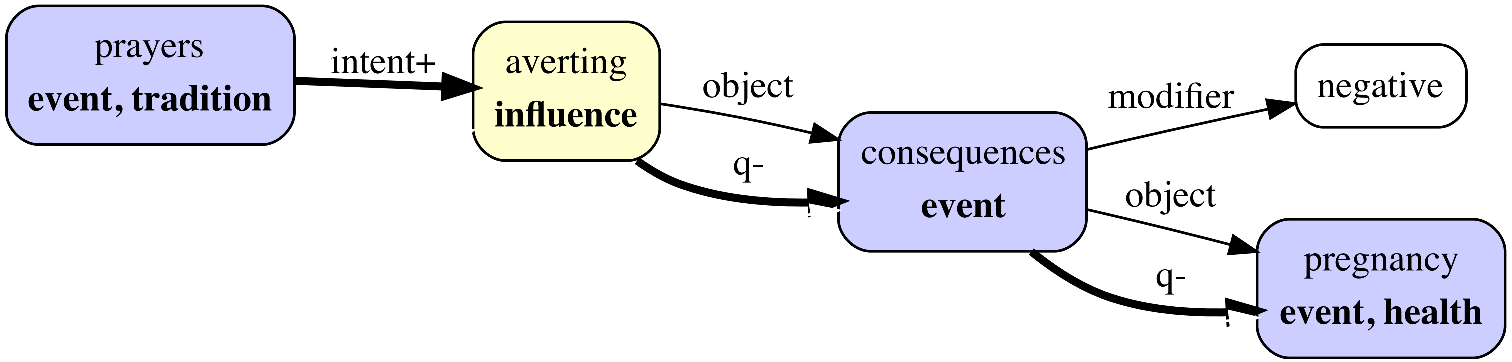

The prayers were aimed at averting negative consequences on the pregnancy.

Propagate Valence via Qualitative Causality

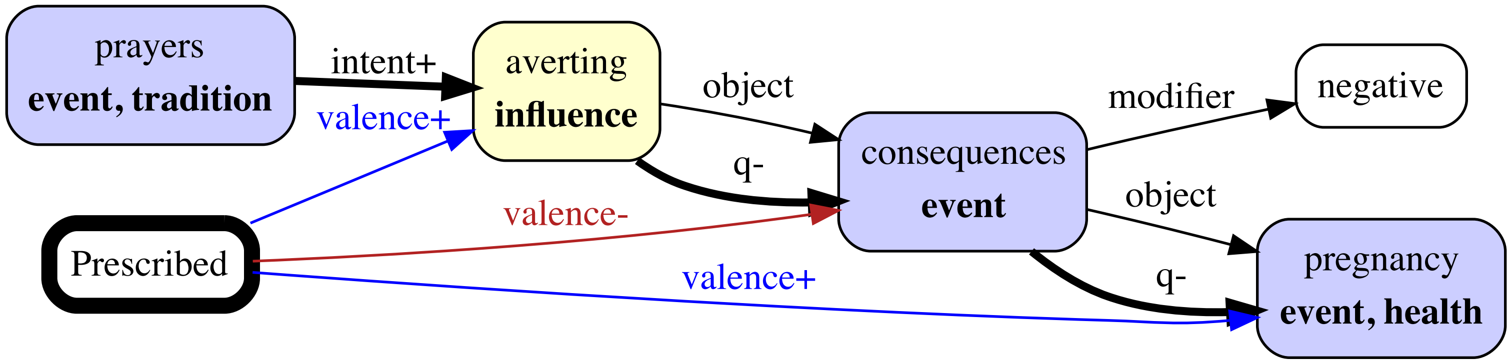

The prayers were aimed at averting negative consequences on the pregnancy.

Generally Prescribed:

- Averting negative consequences on pregnancy is **val+**.
- Negative consequences on pregnancy is **val-**.
- Pregnancy is **val+**.

Propagate Valence via Qualitative Causality

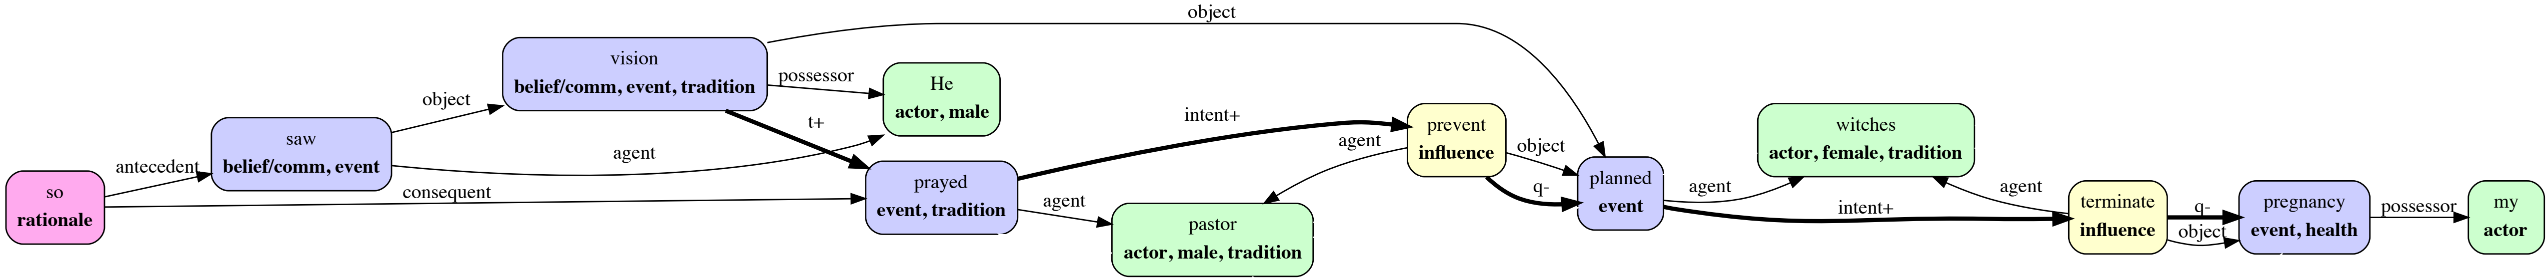

He saw in a vision that the witches had planned to terminate my pregnancy so the pastor prayed to prevent it.

Propagate Valence via Qualitative Causality

The Witches:

- The witches' plan is **val+**.
- Terminating my pregnancy is **val+**.
- My pregnancy is **val-**.

The Pastor:

- Praying to prevent the witches' plan is **val+**.
- The witches' plan is **val-**.
- Terminating my pregnancy is **val-**.
- My pregnancy is **val+**.

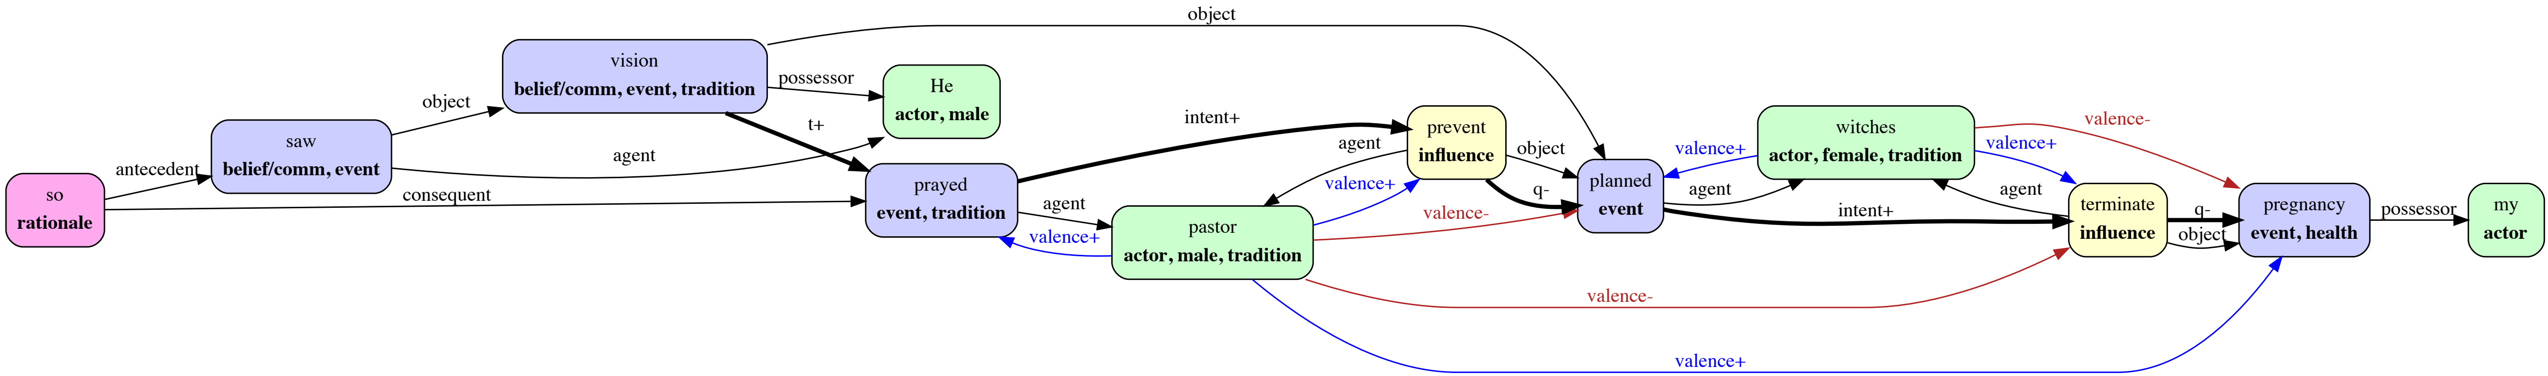

He saw in a vision that the witches had planned to terminate my pregnancy so the pastor prayed to prevent it.

At the example-level: Infer actors' intentions, and quantities they want to minimize, maximize.
At the corpus-level: Infer inter-actor adversity, norms, summaries of heterogeneous local values.

Summary

(Parsed with the scientific claim SpEAR model)

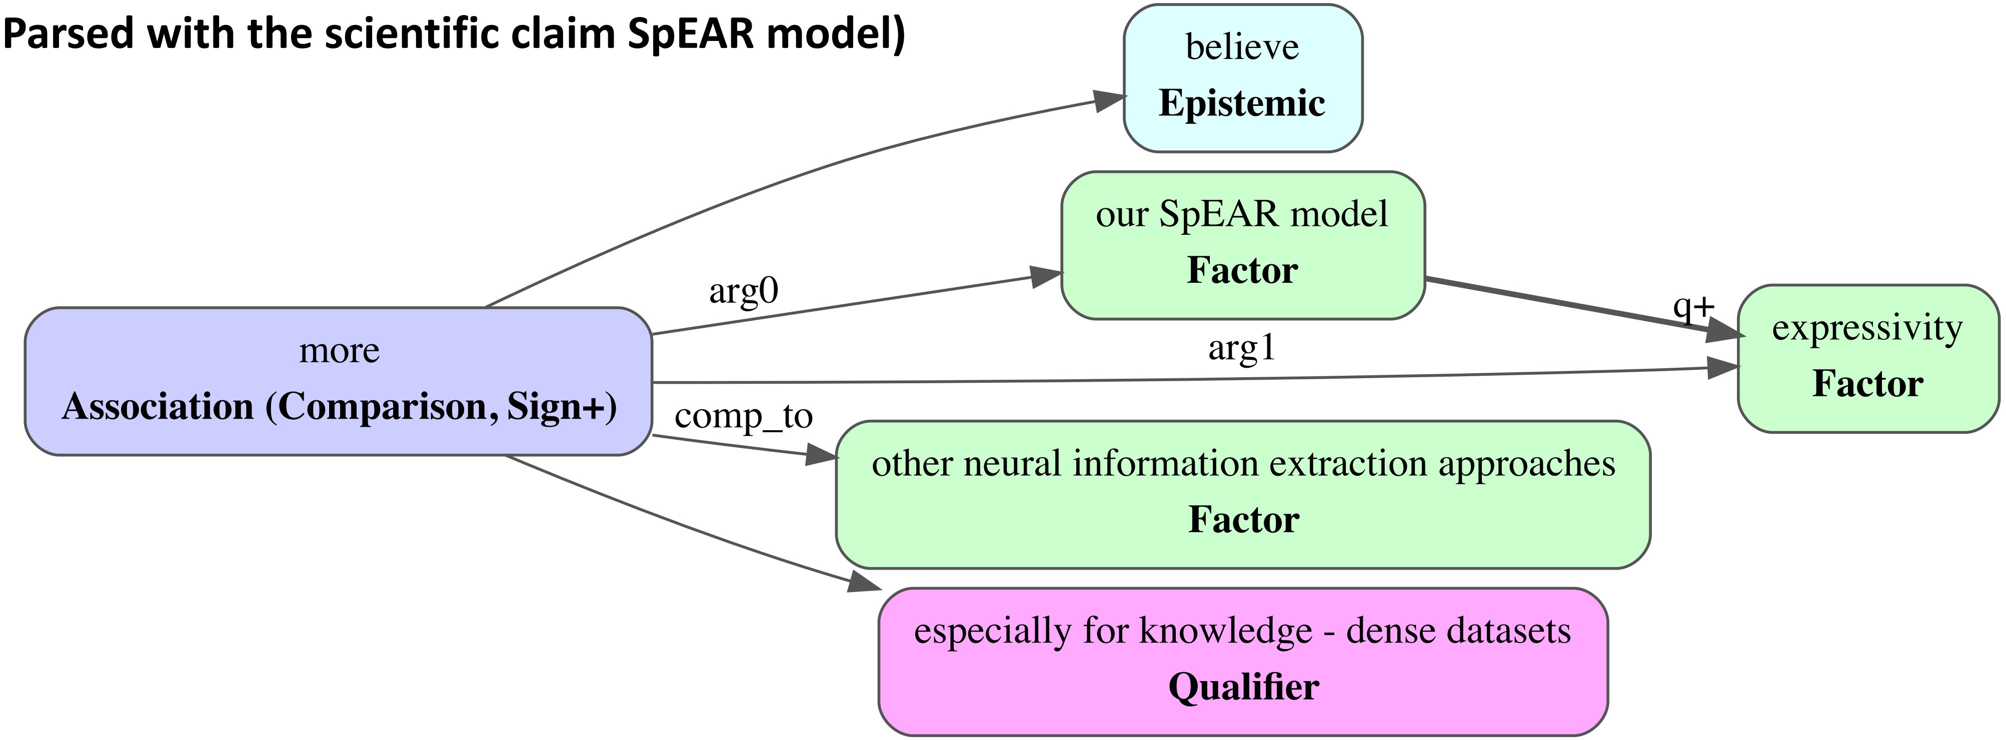

We believe our SpEAR model permits more expressivity than other neural information extraction approaches, especially for knowledge-dense datasets.

Summary

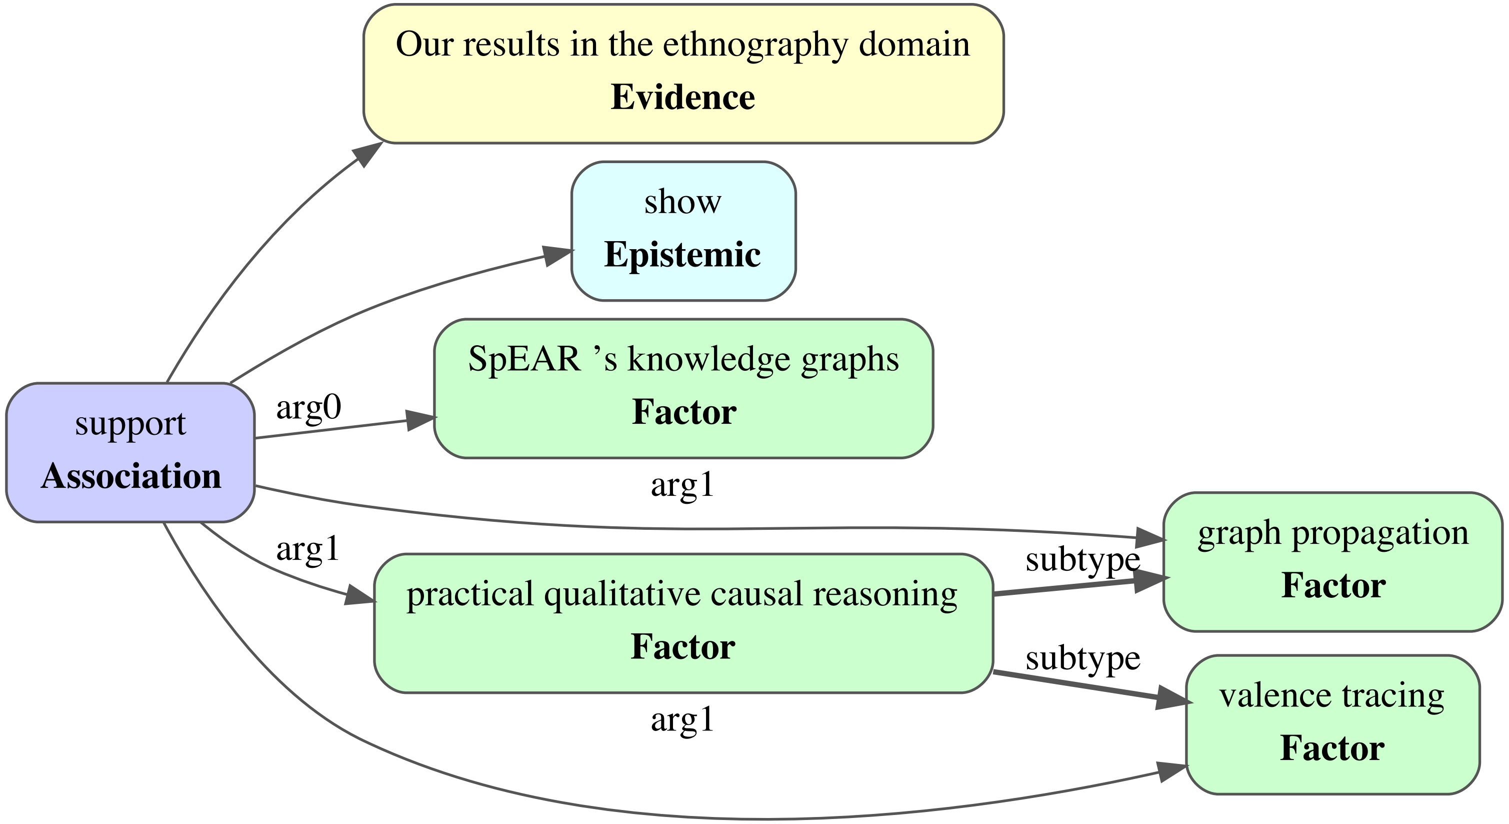

Our results in the ethnography domain show that SpEAR's knowledge graphs support practical qualitative causal reasoning, e.g., graph propagation and valence tracing.

Thank You

(Parsed with the ethnographic SpEAR model)

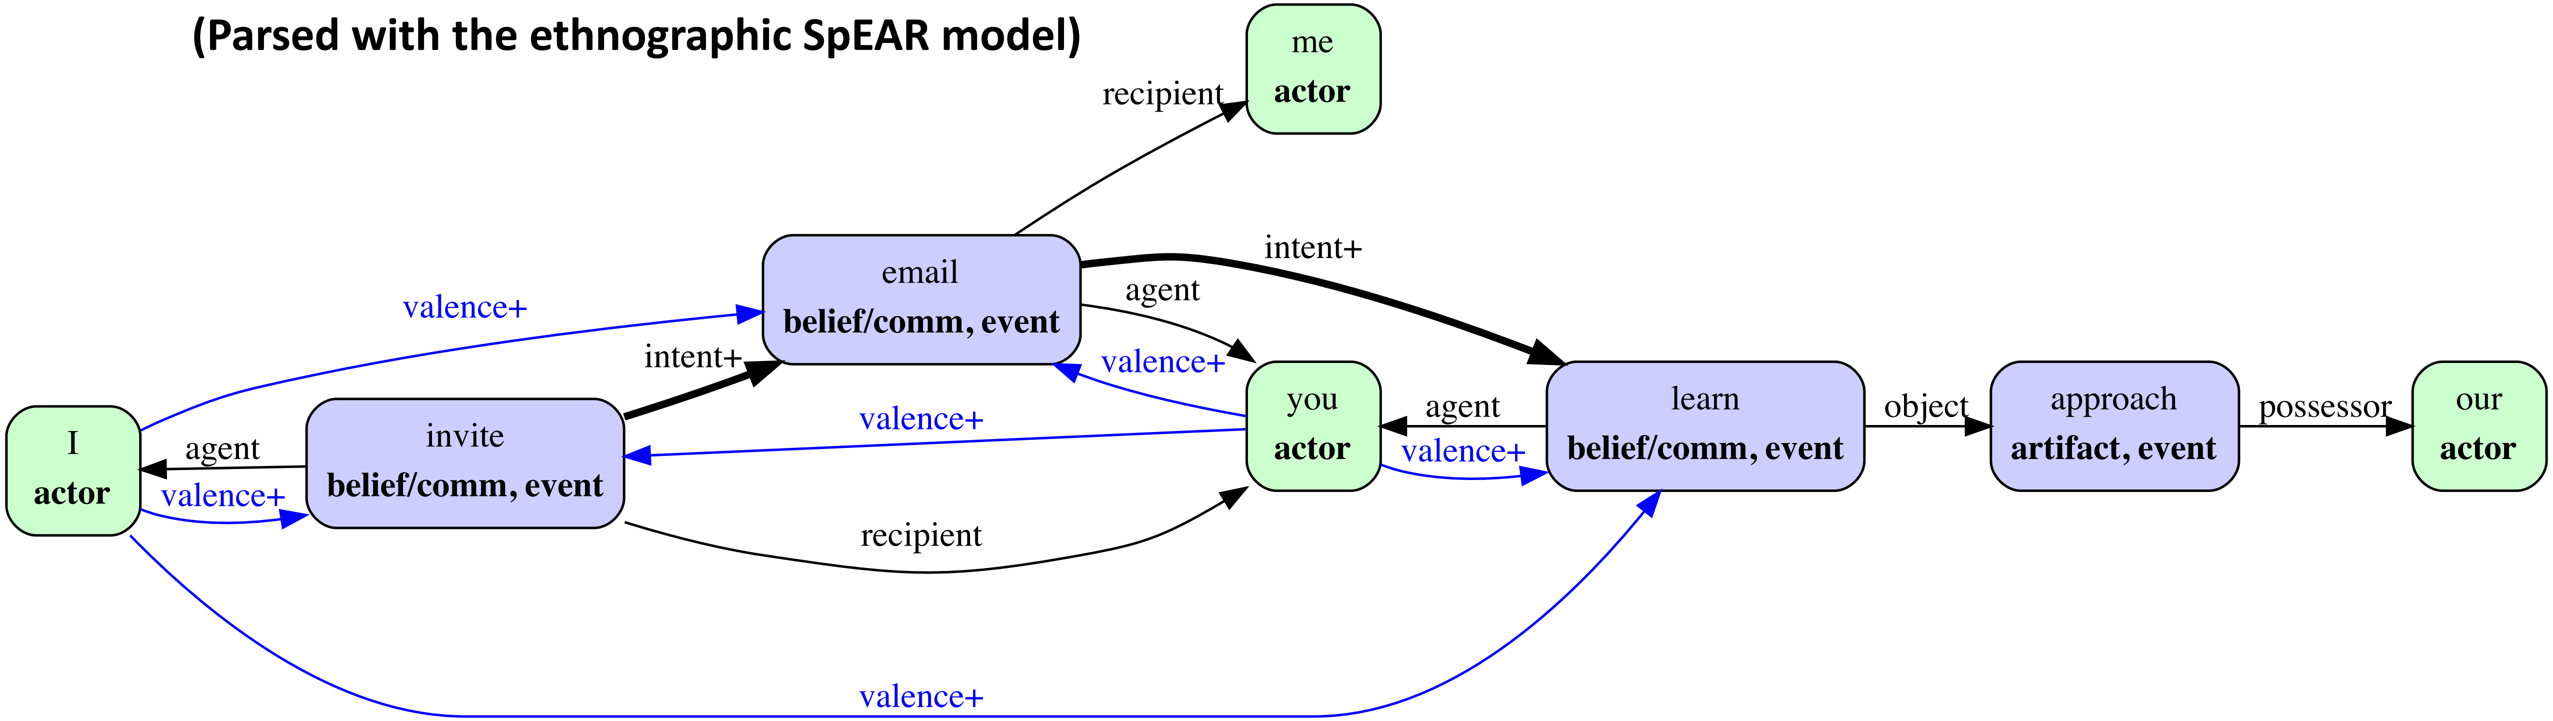

I invite you to email me to learn about our approach.

Scott Friedman (friedman@sift.net),

Ian Magnusson, Vasanth Sarathy, Sonja Schmer-Galunder