Causal schema induction for knowledge discovery

Michael Regan[†♭], **Jena D. Hwang**[‡], **Keisuke Sakaguchi**[+], **James Pustejovsky**[*]

[†] Paul G. Allen School of Computer Science & Engineering, WA, USA

[♭] University of Colorado Boulder, CO, USA

[‡] Allen Institute for AI, WA, USA

[*] Brandeis University, MA, USA

[+] Tohoku University, Sendai, Japan

mregan@cs.washington.edu

Abstract

Making sense of familiar yet new situations typically involves making generalizations about *causal schemas*, stories that help humans reason about event sequences. Reasoning about events includes identifying cause and effect relations shared across event instances, a process we refer to as *causal schema induction*. Statistical schema induction systems may leverage structural knowledge encoded in discourse or the causal graphs associated with event meaning, however resources to study such causal structure are few in number and limited in size.

In this work, we investigate how to apply schema induction models to the task of *knowledge discovery* for enhanced search of English-language news texts. To tackle the problem of data scarcity, we present TORQUESTRA, a manually curated dataset of text-graph-schema units integrating temporal, event, and causal structures. We benchmark our dataset on three knowledge discovery tasks, building and evaluating models for each. Results show that systems that harness causal structure are effective at identifying texts sharing similar causal meaning components rather than relying on lexical cues alone. We make our dataset and models available for research purposes.

1 Introduction

Humans use language to understand stories describing participant interactions in events unfolding over time. To explain novel events in terms of previous experiences, humans rely heavily on *causal schemas*: stories about cause and effect relations that make memory and cognition more efficient (Tversky and Kahneman, 1982; Kahneman, 2012). If such schemas or stories form the basis of human reasoning, perhaps AI systems may similarly learn, store, and manipulate knowledge of causal

(**text**) Karzai said he discussed the issue of **civilian casualties** when he held a **meeting** on **security** and **reconstruction**.

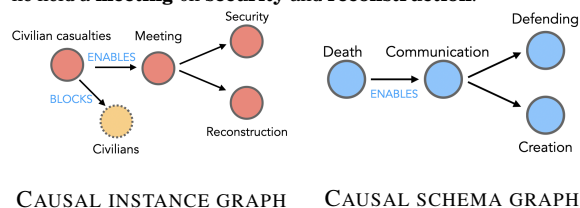CAUSAL INSTANCE GRAPH   CAUSAL SCHEMA GRAPH

Figure 1: A causal schema is either an instance (left) tied directly to a text (top) or a schema graph (right) composed of event types. Edges indicate relations (not all shown) for causation of action and rest (ENABLES and BLOCKS). Graphs include participants, e.g., civilians (dotted orange node, left).

structure for model interpretability or reasoning applications. However, datasets to support studies of causal schemas with natural language processing (NLP) methods are few and far between, a problem we set out to address in this work.

Making a dataset for the computational modeling of causal relations described in language is challenging, and so most existing resources are limited in size and focus on explicit causality at the sentence level. We introduce TORQUESTRA, a dataset of implicit and explicit causal relations at the discourse level to support language studies using statistical methods (e.g., large language models). Our premise is that for human interpretability, causal stories are best represented as graphs, opening up decades of formidable research in graph theory that we can apply to our tasks.

In Fig. 1, we show a pair of TORQUESTRA directed graphs. An instance graph (left) represents the causal story associated with a single text with short descriptions of events and participants as nodes. The corresponding schema (right) is a generalization of this causal story with event types for nodes, giving a means of inferring how different event instances may be similar in predictable ways.

Results of knowledge discovery experiments using TORQUESTRA demonstrate that graph-based methods help identify texts that describe event sequences sharing similar causal structures as well as lexical features, with performance in clustering and schema matching experiments comparable to strong baselines that rely on lexical patterns alone. Through our experiments, we highlight the versatility of the dataset, with the hope of encouraging future research into causation and schemas in NLP.

As we study the inference of latent causal stories given textual descriptions of event sequences, our dataset, TORQUESTRA, may help answer questions such as: (1) In what ways are temporal, causal, event hierarchical, and schema structures related? And, (2) How well do statistical methods such as pre-trained language models help with tasks that resemble causal reasoning? To address these questions, our contributions include:

- (**Theoretical**) We study participant-centered causal structure, a relatively unexplored approach to discourse modeling, for which we define fine-grained causal relations based on physical models of causation;

- (**Dataset**) We present a dataset to analyze the temporal, event, schema, and causal structures described in natural language text; and,

- (**Empirical**) We carry out experiments in structured generation for knowledge discovery, testing the suitability of a general purpose commonsense model distilled on symbolic knowledge for our data and tasks.

We first explore background in schema research (§2) and define causal structure from multiple perspectives including our own (§3). We then take a close look at our dataset, TORQUESTRA, including details about annotation and evaluation (§4). To demonstrate the versatility of the dataset, our experiments include: causal instance graph generation, causal graph clustering, and causal schema matching, and we design and build models and metrics for each (§5). We report baseline results with large language models and graph neural networks (§6), concluding with remarks on challenges and opportunities of schema understanding research.

2 Causal schemas

Schemas, cf. *scripts* and *frames*, are high-level semantic structures for event sequences such as going to restaurants, crime investigations, and investing money (Minsky, 1974; Fillmore, 1976; Schank and Abelson, 1977), a coherent story or pattern of interactions distinct in memory.

2.1 Why schemas?

Schemas help us reconstruct, order, and make predictions about events, about events' relative *salience* and about the *centrality* of event participants. Cognitive processes such as generalization, induction, and intuitive notions of physics and psychology (Talmy, 1988; Tenenbaum et al., 2011) are associated with causal cues encoded in language (Croft, 2012). Together, these base elements give rise to *causal reasoning*, a defining feature of human cognition and possibly one day of AI systems as well (Lake et al., 2017; Schölkopf et al., 2021).

2.2 Causal schema induction

In AI, semantic understanding or analysis is viewed as "abduction to the best explanation" (Hobbs et al., 1993). Abductive reasoning is tightly associated with *induction*, which we view as abduction to the best *high-level* explanation. The causal schema induction task is: given a text, infer high-level semantics for an event sequence using explicit (textual) and implicit (commonsense) knowledge. Consider the example events 1a and 2a.

(1)   a. I passed the salt to you.
      b. TRANSFER $\nrightarrow$ CHANGE OF OWNERSHIP

(2)   a. I passed the money to you.
      b. TRANSFER $\rightarrow$ CHANGE OF OWNERSHIP

We formalize our notion of schemas using event types (1b and 2b) from FrameNet (Fillmore et al., 2003), with arrows denoting causal relations, either lack of enablement ($\nrightarrow$) or enablement ($\rightarrow$). This formalization helps represent typical human experiences, e.g., not all TRANSFER events enable (or *imply*, *entail* or *cause*) a CHANGE OF OWNERSHIP.

Human understanding of which logical conclusions are appropriate in a given context is integral to causal commonsense reasoning. In this paper, we examine how well large language models, e.g., GPT2/3 (Radford et al., 2019; Brown et al., 2020), perform related schema induction tasks.

2.3 Related work

In this section, we briefly describe relevant background in research on stories, event temporality, schemas, semantic search, and challenges of collecting causal data from natural language texts.

Stories and time. A *story* is a temporal ordering of events (Labov and Waletzky, 1967) characterized by change of state and participant interaction (Croft, 2012; Croft et al., 2017). In NLP, research into story understanding has emerged from studies of temporal relations (Allen, 1983; Mani et al., 2006) using temporal data for model development and evaluation (Pustejovsky et al., 2003). Temporal event meaning is nuanced, evident in work on multiple meaning axes (Ning et al., 2018b), temporal aspect (Donatelli et al., 2018), and the relative duration of events (Zhou et al., 2021).

Schemas as temporal structures.   Human knowledge is encoded in stories as *schemas* (Schank and Abelson, 1995), prototypical event sequences for common situations. In NLP, schemas are temporal structures, e.g., narrative event chains (Chambers and Jurafsky, 2008), for event schema induction (Chambers, 2013), timeline construction (Wen et al., 2021), temporal schema induction (Li et al., 2021), future event prediction (Li et al., 2022), and partially-ordered temporal schema generation (Sakaguchi et al., 2021). However, temporal knowledge is complex and inherently noisy (Ning et al., 2018b), likely limiting advances in automated schema understanding systems.

Temporal and causal structures are related. Some lines of work integrate both temporal and causal perspectives, including narrative storylines (Caselli and Vossen, 2016) and representations for temporal and causal networks (Bethard et al., 2008; Berant et al., 2014; Mirza and Tonelli, 2016; O'Gorman et al., 2018). However much of this research does not directly address schemas, which we consider crucial for improved AI reasoning about stories, at the very least in an evaluation context.

Temporal and causal datasets. Our dataset is similar to efforts to crowdsource plot graphs (Li et al., 2013), collect graphical schemas for everyday activities (Sakaguchi et al., 2021), and apply text-graph pairs for temporal reasoning (Madaan and Yang, 2021). Our work differs in that we integrate knowledge of causal, temporal, event, and schema structures in a single dataset.

Semantic search. Knowledge discovery can be framed as semantic search: identifying texts that share semantic structure. Heavy lifters in information retrieval are methods like BM25 and TF-IDF (sparse retrievers), often combined with text embedding similarity metrics (dense retrievers) (Chen et al., 2022). In work close to ours, similarity can also be measured using sentence meaning representations (Bonial et al., 2020), which we extend to study causal structure at the discourse level.

Challenges in making NLP causal datasets. Due to the complexities of faithfully assessing causal relations, natural language datasets have focused mostly on explicit causal markers (Mirza and Tonelli, 2014; Dunietz et al., 2017) typically at the sentence level (Tan et al., 2022). In contrast, we seek to identify implicit (commonsense) and explicit causal relations at the discourse level, resulting in a dataset with at least 6x more causal relations than other resources, as we see in Table 1.

Dataset	# docs	# causal rels
EventCausality (Do et al., 2011)	25	580
Causal-TB (Mirza and Tonelli, 2014)	183	318
RED (O'Gorman et al., 2016)	90	1000
Ning et al (2018a)	25	172
ESTER (Han et al., 2021)	2000	4k
Causal News Corpus (Tan et al., 2022)	3.5k sentences	1600 (train)
Ours (TORQUESTRA$_{human}$)	3k texts	24k
Ours (TORQUESTRA$_{auto}$)	6k long texts	75k

Table 1: Existing human-curated datasets of causal relations in written text are relatively limited in size. Our base dataset (2nd from bottom) is at least 6x greater in number of causal relations compared to previous work.

3 Causal story framework

Causality is complex, as centuries of research stand to remind us. For one, commonsense causal reasoning goes beyond mere notions of necessary and sufficient conditions (Minsky, 1974; Hobbs, 2005). For another, causal viewpoints depend on perspective. Which causal dimensions of event sequences are humans most likely to agree on? To find answers to this question, we examine viewpoints from physics, neuroscience, philosophy, epidemiology, and cognitive semantics.

3.1 Defining causal relations

In Newton's law of inertia (Newton, 1687), cause and effect relations are viewed in terms of **causation of action** (an external force puts an object into motion) and the **causation of rest** (an external force brings an object to rest). We argue that physical models for the acceleration and deceleration of an object through the addition of energy correspond closely with factors leading to the starting and ending of events, drawing on work in cognitive semantics (Talmy, 1988; Croft, 2012) and psychology (Wolff, 2007) where causal relations are conceived of as tendency to action and rest, force, opposition to force, and the overcoming of force.

Rel	Sub-relation	Description	Verbs/concepts (exs.)	Example
ENABLES	BEGINS	Prototypical causation of action	cause, start	Oleg started the ball rolling.
	ADDS	Acceleration; cf. sufficient condition	contribute, help	Olga kept the ball rolling.
	ALLOWS/LETS ACTION	Inaction allows action	let, allow, permit	Olaf let the ball roll (by not acting to stop it).
	PREVENTS REST	Remove barrier so action can continue	free, maintain	Oksana removed obstacles to the ball rolling.
	WITHOUT EFFECT	Despite expectations, no enabling effect	despite, even though	Despite our efforts, we couldn't get the ball rolling.
	UNKNOWN	Uncertainty of enabling relation	questions, modality	Did anybody/anything start the ball rolling?
BLOCKS	ENDS	Prototypical causation of rest	stop	Oleg stopped the ball rolling.
	DISRUPTS	Reduction of momentum	hinder, resist, slow	Olga slowed the ball down.
	ALLOWS/LETS REST	Inaction leads to rest	not help	Olaf let the ball stop rolling.
	PREVENTS ACTION	Barrier to action	refrain, forbid, hold	Oksana prevented the ball from rolling.
	WITHOUT EFFECT	Despite expectations, no blocking effect	despite, even though	We tried but could not stop the ball.
	UNKNOWN	Uncertainty of blocking relation	questions, modality	Did anybody/anything stop the ball?

Table 2: Causal relations reflect the dual concepts of ENABLES (≈makes more likely), shorthand for *causation of action*, and its counterpart BLOCKS (≈makes less likely), shorthand for *causation of rest*. More fine-grained sub-relations (second column) are symmetric, e.g., the most prototypical causal relation ENABLES-BEGINS corresponds to BLOCKS-ENDS, ENABLES-ADDS corresponds to BLOCKS-DISRUPTS, etc. The sub-relation WITHOUT EFFECT denotes the absence of expected causality for events that happen or do not happen *despite expectations*, a challenging task for machines and often overlooked in other datasets.

Likewise, in neuroscience physical causal mechanisms are related to excitatory and inhibitory synapses making neuron responses either more or less likely (Purves et al., 2017), similar to views in scientific philosophy (Reichenbach, 1956) and causation in epidemiology (Parascandola and Weed, 2001). Across views, causal factors are associated with that which increases or decreases the likelihood of events, notions we integrate into definitions for causal structure we present in Table 2.

3.2 Participant-centered causal structure

Modeling causal relations depends not only on how one event (a precondition or *cause*) is related to a subsequent event (a postcondition or *effect*), but also to the direct role of event *participants*, often the grammatical subjects and objects associated with event structure (Talmy, 1988; Croft, 2012). Participant-centeredness is featured in studies of narrative (Propp, 1968; Caselli and Vossen, 2016; Brahman et al., 2021), of participant states (Ghosh et al., 2022; Vallurupalli et al., 2022), and of disease where organisms are conceived of as causative agents, e.g., the pathogen *tubercle bacillus* CAUSES *tuberculosis*.

In the participant-centered graph in Fig. 2 we show how participants and events causally interact. Typical causal relations are between events, e.g., the *ousting of a leader* may end or BLOCK a *conflict*. In a participant-centered approach, people and things directly act on one another and also act as the initiating agents or causal endpoints of events, e.g., the *rebels* BLOCK the *leader* who in turn ENABLES the conflict, etc.
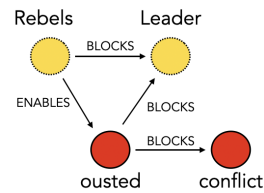

Figure 2: An example participant-centered instance graph for "The rebels ousted the leader to end the conflict." Causal graphs consist of two types of nodes: participants (top, orange) and events (bottom, red). Here, the rebels enable the ousting event, thus blocking the leader and the conflict as well.

4 TORQUESTRA

TORQUESTRA is a *causal schema library*: a dataset[1] of text-graph-schema units for research in schema induction and, more broadly, knowledge discovery (see §5). At its core, a TORQUESTRA data instance is a newswire text paired with causal (G_{causal}), temporal (G_{temp}) and event (G_{event}) structures. For notation, see Table 3.

Symbol	Meaning
$G_{causal} \in \mathcal{G}$	Instance causal graph assoc. w/ text
$S_{causal} \in \mathcal{S}$	Schema causal graph assoc. w/ text(s)
$p \in \mathcal{P}$	Participant node in a causal graph
$\langle p_i, rel, p_j \rangle$	Causal relation (see Table 2) between participants p_i and p_j
G_{temp}	Temporal graph (Ning et al., 2020)
G_{event}	Event graph (Han et al., 2021)
V_{maven}	Node w/ event types (Wang et al., 2020)
ϕ_{maven}	Set of hierarchical event types

Table 3: Notation used in this paper.

In this section, we briefly describe the resources we used to make TORQUESTRA, including notes about texts, size, graphs, and data annotation.

---

[1] https://github.com/fd-semantics/causal-schema-public

4.1 Temporal and event structure knowledge

TORQUE + ESTER = TORQUESTRA. We wanted a dataset for a joint study of temporal, causal, and event structures. To this end, we examined question-answer (QA) datasets for temporal relations TORQUE (Ning et al., 2020) and event structures ESTER (Han et al., 2021). After noting both drew texts from TempEval3, a subsequent analysis revealed the datasets shared 700 text snippets, an intersection of data that we used to form core TORQUESTRA, as illustrated in Fig. 3.
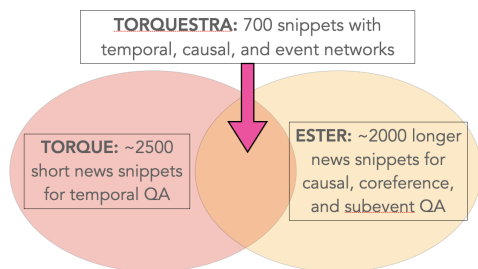

Figure 3: The core of TORQUESTRA is drawn from texts with rich QA annotations from two existing resources: TORQUE (Ning et al., 2020) and ESTER (Han et al., 2021).

4.2 Texts

Texts in TORQUESTRA are English-language newswire snippets from TempEval3 (UzZaman et al., 2013) and Wikipedia. The texts cover a number of typical news domains, including politics, sports, and business. Texts are mostly multiple sentences (98%+), with mean text length between 60-300 subword tokens (Byte Pair Encoding (Sennrich et al., 2016)).

4.3 Dataset size

Our manually constructed dataset consists of three slices of data: TORQUE (2500 exs), ESTER (700 exs), and WIKI-CRIME (200 exs), each aligned with up to four semantic networks. For details about data slices, see Appendix A.2, Fig. 6.

4.4 Causal instance graphs

TORQUESTRA consists of causal instance graphs for events described in text (see Fig. 2). A causal graph $G_{causal} = (V, E)$ is directed and at times cyclic, with vertices V for salient events and participants and edges E being causal relations. Nodes in the graphs are natural language descriptions for events and event participants, typically of subject-verb-object form. For graph statistics, see Appendix A.2, Fig. 7.

4.5 Causal schema graphs

As counterparts to G_{causal}, schema causal graphs S_{causal} are generalizations for event sequences using event types from $V_{maven} \subset \phi_{maven}$[2]. Annotators also add free-form event labels for cases not represented in ϕ_{maven}, e.g., event types such as INTERNATIONAL_RELATIONS, amounting to >400 event types observed (with details in Appendix A.3).

Compare the instance graph from earlier (Fig. 2) to the schema graph in Fig. 4: the 'ousting of the leader' blocks the 'conflict' in the instance graph, which in the schema graph is generalized to a CHANGE OF LEADERSHIP that blocks a MILITARY OPERATION.
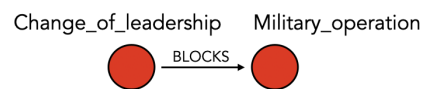

Figure 4: A causal schema graph for the event sequence "The ousting of the leader ended the conflict" using frame semantic labels as nodes.

Ontological questions for causal schemas remain open. In TORQUESTRA, causal schema are associated with event frame semantics at the node- or subgraph-level, one step of the data collection process that we discuss next.

4.6 Data annotation

The dataset is compiled through manual and automated means, subsets of data we refer to as TORQUESTRA$_{human}$ and TORQUESTRA$_{auto}$. In this subsection, we focus on manual annotation efforts.

TORQUESTRA$_{human}$ consists of approximately 30K spans of text corresponding to graph nodes and 48K additional labels for nodes and edges. Annotation consisted of four main tasks: Given a short text and commonsense knowledge, identify and label causal participants (nodes), event types (for nodes/graph), causal relations (edges), and salient causal chains.

For annotation, we relied on a group of eight (8) in-house undergraduate and graduate students with backgrounds in linguistics and computer science which we found could faithfully recreate the causal graphs we envisioned. Core TORQUESTRA$_{human}$ required approx 250 hours with annotators earning between $16-25/hr. For more details about the annotation process, guidelines, evaluation, and prompt engineering, see Appendix A.4.

---

[2] $|\phi_{maven}| = 168$ FrameNet event types

5 Tasks and experimental methods

TORQUESTRA supports the induction and knowledge discovery tasks illustrated in Fig. 5: (1) *causal instance graph generation*, (2) unsupervised *causal graph clustering*, and (3) *causal schema matching*. We briefly describe each in turn.
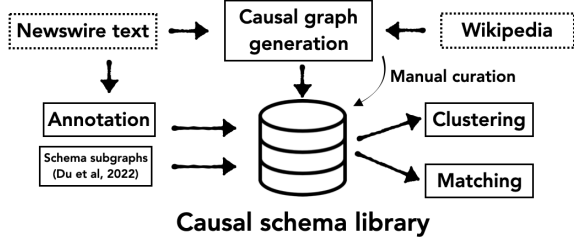

Figure 5: Data-task pipeline. With texts from different sources (dotted boxes), we annotate, generate, and collect causal stories for a *schema library*, a repository we use for *causal graph clustering* and *schema matching*.

5.1 Causal instance graph generation

The first task is *causal instance graph generation*. As in work on narrative planning (Riedl and Young, 2010), learning mini knowledge graphs as world models (Ammanabrolu and Riedl, 2021), and temporal graph generation (Madaan and Yang, 2021), we generate graphs conditioned on text and, in an extension to previous work, also condition on event semantics (e.g., temporal structures). We compare few-shot GPT-3 (Brown et al., 2020) with fine-tuned GPT2-XL and knowledge distilled GPT2-XL$_{distill}$(West et al., 2022) (60/40 train/dev split).

Generation evaluation. Perplexity and n-gram overlap metrics such as METEOR (Banerjee and Lavie, 2005) are of limited use as proxy measures for the faithfulness and coherence of generated causal stories. So, we also manually evaluate triples, reporting *correctness* (% of accurate edges) and *completeness* (% of causal graph generated).

5.2 Causal graph clustering

The second task is *causal graph clustering*, unsupervised clustering of schema instance graphs (TORQUESTRA$_{auto}$). The objective of this task is to study the effectiveness of similarity metrics for texts using lexical features and graph embeddings.

Data. As out-of-domain data for testing, we used MAVEN (MAssive eVENt detection) (Wang et al., 2020), a collection of 3K+ Wikipedia articles that targets open domain event understanding systems, adopting 168 labels from FrameNet (Fillmore et al., 2003) organized in an event hierarchy.

Models. For lexical similarity baselines, we use standard implementations of tf-idf[3] for sparse vectors and a SentenceTransformer (Reimers and Gurevych, 2019) for dense text embeddings.

For graph embeddings, we first encode graph nodes using DeBERTa (900M model) (He et al., 2021), and assign scalar values to edges (+1 for ENABLES and −1 for BLOCKS). We then train a graph attention network (Veličković et al., 2018) via self-supervision masking random nodes. For further comparison, we also compose graph embeddings using the FEATHER algorithm (Rozemberczki and Sarkar, 2020) based on random walks.

We cluster embeddings using standard K-means[4] with k=6 for the number of clusters and fix the number of observations to n=25 for evaluation. We then examine results, measuring similarity using automated metrics we outline next.

Clustering metrics. We report purity, adjusted Rand index (ARI), and V-measure (VM) (Rosenberg and Hirschberg, 2007) using as ground truth topic labels for each article (Wang et al., 2020) mapped to a smaller ontology[5].

To measure similarity accounting for multiple labels, we propose a new metric: *event cluster purity*, estimating for each cluster a true label e_j as the top-j event types observed[6]. In these experiments, we look at the top-10 event types observed, j=10, so e_{10}^m denotes the ground truth event type vector for cluster m. We then compare this ground truth with a human-annotated k-hot event vector for each graph, e^G. For N clusters M, the metric is defined:

$$\text{purity}_{event} = \frac{1}{N} \sum_{m \in M} \sum_{e^G \in m} \max_{e_{10}^m \in \mathbf{E}} |e_{10}^m \cap e^G| \quad (1)$$

5.3 Causal schema matching

As a variant of exemplar matching, the aim of causal schema matching is to identify induced MAVEN graphs (TORQUESTRA$_{auto}$) most similar to curated schemas from two sources: core TORQUESTRA$_{human}$ (short texts) and an existing schema library (Du et al., 2022). In the latter case, we match to schema chapters (individual subgraphs) for ease of evaluation.

Previous work has investigated methods to align schema nodes (Du et al., 2022), and methods for subgraph matching exist (Rex et al., 2020).

---

[3]https://en.wikipedia.org/wiki/Tf-idf
[4]https://scikit-learn.org/stable/modules/clustering.html
[5]e.g., 'hurricane' and 'earthquake' are both DISASTER
[6]Reminiscent of Jaccard similarity

Experiment	Model	Output | Input	Metrics			
			METEOR	Correct	Complete	# triples eval
Graph generation 1	GPT-3 (7-shot)	$p(G_{causal}\|text, G_{temp})$	0.28	0.50	0.33	224
(validation set)		$p(G_{causal}\|text, G_{event})$	0.26	0.55	0.25	180
(supervised, 60/40 split)	GPT2-XL	$p(G_{causal}\|text, G_{temp})$	0.27	0.52	0.23	180
	GPT2-XL$_{distill}$	$p(G_{causal}\|text)$	0.27	0.58	0.35	120
		$p(G_{causal}\|text, G_{temp})$	0.29	0.60	0.38	120
		$p(G_{causal}\|text, V_{event})$	0.34	0.60	0.40	300
		$p(G_{causal}\|text, G_{temp}, V_{event})$	**0.41**	**0.65**	**0.42**	360
Graph generation 2	GPT2-XL$_{distill}$	$p(G_{causal}\|text, G_{temp})$	n/a	0.56	0.33	320
(test set)		$p(G_{causal}\|text, V_{event})$	n/a	0.59	0.36	360

Table 4: Results for causal graph (G_{causal}) generation using GPT-3 (Brown et al., 2020), GPT2-XL (Radford et al., 2019) and GPT2-XL$_{distill}$ (West et al., 2022). We condition on texts + temporal networks (G_{temp}), event structure (G_{event}), and hierarchical events (V_{event}), automatically evaluating with METEOR and manually evaluating correctness and completeness. Pairs of results underlined (or dotted) illustrate important points we discuss in §6.

Nonetheless, our experiments show the effectiveness of using graph embedding similarity as a step in identifying relevant schemas given a query.

Using the same models as our clustering experiments, we randomly select 50 Wikipedia articles to match with our schema library (RESIN (Du et al., 2022) + core TORQUESTRA) and examine the top-5 matched schemas using text topic labels, event type overlap, and graph visualizations for qualitative analysis. We report mean average precision (MAP) and mean reciprocal rank (MRR) as the accuracy of the ranking of the most relevant text. For more details about metrics and the evaluation tool, consult the Appendix A.8 and website.

6 Results and discussion

We present results for causal instance graph generation using manual and automatic evaluation in Table 4 and results for causal graph clustering and schema matching in Tables 5 and 6.

We report mean results for a minimum of three different model runs varying random seeds (of graph neural networks) and hyperparameters (# epochs, block size, p-sampling rate, etc.). We do not exhaustively explore settings nor compare with language models outside the GPT family. Experimentation leads to the following observations.

Large language models can generate complex structured representations. Experiments show (Table 4) we can generate symbolic causal knowledge in the form of directed, branching causal graph structures with multiple events and participants. We expect that research into interpretable, neuro-symbolic, stepwise reasoning using generative models may build upon this progress in structure prediction.

Model	Input	Metric			
		purity	ARI	VM	purity$_{event}$
embedding	text	0.96	0.95	0.95	4.36
TF-IDF	text	**0.98**	**0.97**	**0.97**	**5.47**
FEATHER	graph	0.82	0.20	0.33	4.09
GAT	graph	0.83	0.46	0.49	4.69

Table 5: Results for causal graph clustering (higher is better). Evaluation is based on single labels (for first three metrics), with purity$_{event}$ (Eq. 1) based on most frequent event types.

Method	Matching	Metric	
		MAP	MRR
TF-IDF	text-to-text	0.36	0.32
GAT	graph-to-graph	0.48	0.36
TF-IDF	text-to-schema	0.59	0.35
GAT	graph-to-schema	**0.68**	**0.43**

Table 6: First (top), we match TORQUESTRA to Wikipedia texts (MAVEN) using TF-IDF and graphs encoded with a graph attention network (GAT). For schema matching (bottom), we match MAVEN graphs to our causal schema library.

Conditioning on structural knowledge improves generation performance, in some cases. We evaluate if conditioning on temporal, event, and event type networks helps improve generated causal graph correctness and completeness and find that temporal and event structures appended to raw text result in more correct (Table 4, +7%), more complete (+7%) graphs than raw text alone. Overall, semantic signals jointly increase performance over conditioning on raw text alone using validation data (texts 100-150 tokens in length).

More specifically, we experimented with various forms of concatenated text and structures, including text alone, text + G$_{temp}$, and text + G$_{temp}$ + V$_{event}$. With in-distribution data (top half of Table 4), the

most rich input ($G_{temp} + V_{event}$) led to the best generation results. In contrast, with out-of-domain texts (bottom half), text alone works better than conditioning with 'dense' paraphrases of V_{event}, with the length of test texts (2-4x longer than avg. validation) likely a factor.

The student surpasses the teacher. Following work in knowledge distillation (West et al., 2022), experiments show GPT2-XL$_{distill}$ (student model, trained on knowledge graph triples, fine-tuned on TORQUESTRA) outperforms few-shot GPT-3 (teacher model, trained to predict next word) (Table 4, +10%). Further, GPT2-XL$_{distill}$ outperforms original GPT2-XL in correctness (+8%) and completeness (+15%) (input text + G_{temp}), evidence the distilled model learns causal structure, suggesting that we need commonsense knowledge for more complete and correct causal graph generation.

Lexical methods for clustering texts are generally better. Unsurprisingly, tf-idf significantly outperforms graph similarity methods across all metrics (Table 5). We note tf-idf clusters are 'quite' homogeneous, due in part to the provenance of the test data: Wikipedia articles automatically selected and labeled with topics (Wang et al., 2020), likely with similar methods as ours.

Advantages of matching using graph methods versus words alone. We measure similarity of event sequences for clustering comparing text-to-texts and graph-to-graphs, and for matching experiments comparing graph-to-schemas and text-to-schemas. Advantages of our system are evident matching graph-to-schemas (Table 6, +9%).

We find graph-based methods help identify articles with similar causal stories, e.g., graphs with a 4-nary node 'military operation'. Graph-based methods provide a means of measuring conceptual similarities between the causal stories associated with events that may not otherwise be matched. For example, our algorithm finds a high similarity between the 1939 'Invasion of Poland' and a head-on train collision, where both stories involve opposing forces running into each other explosively, with similar (predictable) tragic consequences.

Smaller block size helps identify salient subgraphs. In training, setting block size (the length of input presented to the model) to shorter lengths (e.g., < 300 subword units) provides the model with only a subset of causal triples for each text. As we topologically sort input graphs using breadth-first search, the model learns to generate salient and connected edges. We leave for future work more rigorous evaluation of salience detection using manual annotations we include as part of our data release.

Evaluation is challenging. There are many challenges associated with the evaluation of schema induction systems. On the one hand, lexical overlap and shared entities make two texts similar. On the other, similar causal structures, i.e. the causal schemas that stories share, can be discovered and compared. Still, the weighing of multiple semantic signals remains subjective.

We experimented with various means of evaluation: precision of topic labels (e.g., *man-made disaster*), overlap of event types (text as bag-of-events) and subsets of frequent event types (Eq. 1). We qualitatively assess graph structural similarity, with an automated tool a work in-progress.

We find that multiple measures of schema similarity to be more robust than using a single method, though we also recognize that more work, both theoretical and computational, needs to be done to develop still more reliable tools.

Schema meaning. Previous work views schemas as linear event orderings (Chambers, 2013) and as more complex graph structures (Li et al., 2021; Du et al., 2022). How to further compose atomic meanings into larger semantic units for computational processing remains an open research question. Something like an event ontology of hierarchical event structures likely plays a role in the human conceptualization of event similarity, however, we make no hypotheses about better representations for computational applications.

7 Conclusion

We present TORQUESTRA, a dataset of paired semantic graphs for studies of causal structure at the discourse level. Our experiments in causal graph generation, clustering, and schema matching provide insight into how to leverage TORQUESTRA for knowledge discovery of latent causal structures of news texts, comparable to or outperforming search methods based on lexical similarity alone.

Research in knowledge discovery using causal schema induction will be of interest to historians, journalists, and health researchers looking for new angles on the study of narratives and stories. To support such research, we make our dataset, starter code, and evaluation tools publicly available[7].

---

[7]https://fd-semantics.github.io/

Acknowledgements

References

James F. Allen. 1983. Maintaining knowledge about temporal intervals. *Commun. ACM*, 26(11):832–843.

Prithviraj Ammanabrolu and Mark O. Riedl. 2021. Learning knowledge graph-based world models of textual environments. *arXiv*.

Satanjeev Banerjee and Alon Lavie. 2005. METEOR: An automatic metric for MT evaluation with improved correlation with human judgments. In *Proceedings of the ACL Workshop on Intrinsic and Extrinsic Evaluation Measures for Machine Translation and Summarization*, pages 65–72, Ann Arbor, Michigan. Association for Computational Linguistics.

Jonathan Berant, Vivek Srikumar, Pei-Chun Chen, Brad Huang, Christopher D. Manning, Abby Vander Linden, Brittany Harding, and Peter Clark. 2014. Modeling biological processes for reading comprehension. *EMNLP*.

Steven Bethard, William Corvey, Sara Klingenstein, and James H. Martin. 2008. Building a corpus of temporal-causal structure. In *Proceedings of the Sixth International Conference on Language Resources and Evaluation (LREC'08)*, Marrakech, Morocco. European Language Resources Association (ELRA).

Claire Bonial, Stephanie M. Lukin, David Doughty, Steven Hill, and Clare Voss. 2020. InfoForager: Leveraging semantic search with AMR for COVID-19 research. In *Proceedings of the Second International Workshop on Designing Meaning Representations*, pages 67–77, Barcelona Spain (online). Association for Computational Linguistics.

Faeze Brahman, Meng Huang, Oyvind Tafjord, Chao Zhao, Mrinmaya Sachan, and Snigdha Chaturvedi. 2021. "let your characters tell their story": A dataset for character-centric narrative understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1734–1752, Punta Cana. Association for Computational Linguistics.

Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners. *NeurIPS*.

Tommaso Caselli and Piek Vossen. 2016. The storyline annotation and representation scheme (star): A proposal. *Proceedings of the 2nd Workshop on Computing News Storylines, ACL*, pages 67–72.

Nathanael Chambers. 2013. Event schema induction with a probabilistic entity-driven model. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807, Seattle, Washington, USA. Association for Computational Linguistics.

Nathanael Chambers and Dan Jurafsky. 2008. Unsupervised learning of narrative event chains. In *Proceedings of ACL-08: HLT*, pages 789–797, Columbus, Ohio. Association for Computational Linguistics.

Xilun Chen, Kushal Lakhotia, Barlas Oguz, Anchit Gupta, Patrick Lewis, Stan Peshterliev, Yashar Mehdad, Sonal Gupta, and Wen-tau Yih. 2022. Salient phrase aware dense retrieval: Can a dense retriever imitate a sparse one? In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 250–262, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

William Croft. 2012. *Verbs: Aspect and Causal Structure*. Oxford University Press.

William Croft, Pavlina Peskova, and Michael Regan. 2017. Integrating decompositional event structure into storylines. *Events and Stories in the News Workshop, Association for Computational Linguistics*.

Quang Do, Yee Seng Chan, and Dan Roth. 2011. Minimally supervised event causality identification. In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 294–303, Edinburgh, Scotland, UK. Association for Computational Linguistics.

Lucia Donatelli, Michael Regan, William Croft, and Nathan Schneider. 2018. Annotation of tense and aspect semantics for sentential AMR. In *Proceedings of the Joint Workshop on Linguistic Annotation, Multiword Expressions and Constructions (LAW-MWE-CxG-2018)*, pages 96–108, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Xinya Du, Zixuan Zhang, Sha Li, Ziqi Wang, Pengfei Yu, Hongwei Wang, Tuan Manh Lai, Xudong Lin, Iris Liu, Ben Zhou, Haoyang Wen, Manling Li, Darryl Hannan, Jie Lei, Hyounghun Kim, Rotem Dror, Haoyu Wang, Michael Regan, Qi Zeng, Qing Lyu, Charles Yu, Carl Edwards, Xiaomeng Jin, Yizhu

Jiao, Ghazaleh Kazeminejad, Zhenhailong Wang, Chris Callison-Burch, Carl Vondrick, Mohit Bansal, Dan Roth, Jiawei Han, Shih-Fu Chang, Martha Palmer, and Heng Ji. 2022. Resin-11: Schema-guided event prediction for 11 newsworthy scenarios. In *2022 Annual Conference of the North American Chapter of the Association for Computational Linguistics – System Demonstration Track*.

Jesse Dunietz, Lori Levin, and Jaime Carbonell. 2017. Automatically tagging constructions of causation and their slot-fillers. *Transactions of the Association for Computational Linguistics*, 5:117–133.

Charles J. Fillmore. 1976. Frame semantics and the nature of language. *Annals of the New York Academy of Sciences: Conference on the Origin and Development of Language and Speech*, 280(1):20–32.

Charles J. Fillmore, Christopher R. Johnson, and Miriam R.L. Petruck. 2003. Background to FrameNet. *International Journal of Lexicography*, 16.3:235–250.

Sayontan Ghosh, Mahnaz Koupaee, Isabella Chen, Francis Ferraro, Nathanael Chambers, and Niranjan Balasubramanian. 2022. Pasta: A dataset for modeling participant states in narratives. *arXiv*.

Rujun Han, I-Hung Hsu, Jiao Sun, Julia Baylon, Qiang Ning, Dan Roth, and Nanyun Peng. 2021. ESTER: A machine reading comprehension dataset for reasoning about event semantic relations. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7543–7559, Online and Punta Cana. Association for Computational Linguistics.

Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. 2021. Deberta: Decoding-enhanced bert with disentangled attention. In *International Conference on Learning Representations*.

Jerry R. Hobbs. 2005. Toward a useful concept of causality for lexical semantics. *Journal of Semantics*, 22(2):181–209.

Jerry R. Hobbs, Mark E. Stickel, Douglas E. Appelt, and Paul Martin. 1993. Interpretation as abduction. *Artificial Intelligence*, 63(1):69–142.

Ari Holtzman, Jan Buys, Li Du, Maxwell Forbes, and Yejin Choi. 2019. The curious case of neural text degeneration. *arXiv*.

Daniel Kahneman. 2012. *Thinking, Fast and Slow*. Penguin Books, London.

William Labov and Joshua Waletzky. 1967. Narrative analysis: Oral versions of personal experience. *Journal of Narrative and Life History*, 7(1-4):3–38.

Brenden Lake, Tomer D. Ullman, Joshua B. Tenenbaum, and Samuel J. Gershman. 2017. Building machines that learn and think like people. *Behavioral and Brain Sciences*, 40:1–58.

Boyang Li, Stephen Lee-Urban, George Johnston, and Mark O. Riedl. 2013. Story generation with crowd-sourced plot graphs. In *Proceedings of the Twenty-Seventh AAAI Conference on Artificial Intelligence*, AAAI'13, page 598–604. AAAI Press.

Manling Li, Sha Li, Zhenhailong Wang, Lifu Huang, Kyunghyun Cho, Heng Ji, Jiawei Han, and Clare Voss. 2021. Future is not one-dimensional: Graph modeling based complex event schema induction for event prediction. *EMNLP*.

Manling Li, Ruochen Xu, Shuohang Wang, Luowei Zhou, Xudong Lin, Chenguang Zhu, Michael Zeng, Heng Ji, and Shih-Fu Chang. 2022. Clip-event: Connecting text and images with event structures. *arXiv*.

Aman Madaan and Yiming Yang. 2021. Neural language modeling for contextualized temporal graph generation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 864–881, Online. Association for Computational Linguistics.

Inderjeet Mani, Marc Verhagen, Ben Wellner, Chong Min Lee, and James Pustejovsky. 2006. Machine learning of temporal relations. In *Proceedings of the 21st International Conference on Computational Linguistics and 44th Annual Meeting of the Association for Computational Linguistics*, pages 753–760, Sydney, Australia. Association for Computational Linguistics.

R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. 2002. Network motifs: Simple building blocks of complex networks. *Science*, 298(5594):824–827.

Marvin Minsky. 1974. A framework for representing knowledge. Technical report, Massachusetts Institute of Technology, USA.

Paramita Mirza and Sara Tonelli. 2014. An analysis of causality between events and its relation to temporal information. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2097–2106, Dublin, Ireland. Dublin City University and Association for Computational Linguistics.

Paramita Mirza and Sara Tonelli. 2016. CATENA: CAusal and TEmporal relation extraction from NAtural language texts. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 64–75, Osaka, Japan. The COLING 2016 Organizing Committee.

Isaac Newton. 1687. *Philosophiae Naturalis Principia Mathematica*. Londini, Jussu Societatis Regiæ ac Typis Josephi Streater.

Qiang Ning, Zhili Feng, Hao Wu, and Dan Roth. 2018a. Joint reasoning for temporal and causal relations. In

Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, pages 2278–2288, Melbourne, Australia. Association for Computational Linguistics.

Qiang Ning, Hao Wu, Rujun Han, Nanyun Peng, Matt Gardner, and Dan Roth. 2020. TORQUE: A reading comprehension dataset of temporal ordering questions. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1158–1172, Online. Association for Computational Linguistics.

Qiang Ning, Hao Wu, Haoruo Peng, and Dan Roth. 2018b. Improving temporal relation extraction with a globally acquired statistical resource. *NAACL HLT 2018 - 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies - Proceedings of the Conference*, 1:841–851.

Tim O'Gorman, Michael Regan, Kira Griffitt, Ulf Hermjakob, Kevin Knight, and Martha Palmer. 2018. AMR beyond the sentence: the multi-sentence AMR corpus. In *Proceedings of the 27th International Conference on Computational Linguistics*, pages 3693–3702, Santa Fe, New Mexico, USA. Association for Computational Linguistics.

Tim O'Gorman, Kristin Wright-Bettner, and Martha Palmer. 2016. Richer Event Description: Integrating event coreference with temporal, causal and bridging annotation. *2nd Workshop on Computing News Storylines, ACL*, pages 47–56.

Jessica Ouyang and Kathleen McKeown. 2015. Modeling reportable events as turning points in narrative. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 2149–2158, Lisbon, Portugal. Association for Computational Linguistics.

Mark Parascandola and D L Weed. 2001. Causation in epidemiology. *Journal of Epidemiology & Community Health*, 55(12):905–912.

Vladimir Propp. 1968. *Morphology of the folktale, 2nd ed.* University of Texas Press.

Dale Purves, George J. Augustine, David Fitzpatrick, William C. Hall, Anthony LaMantia, Richard D. Mooney, Michael L. Platt, and Leonard E. White. 2017. *Neuroscience, 6th ed.* Sinauer Associates.

James Pustejovsky, José M. Castaño, Robert Ingria, Roser Saurí, Robert J. Gaizauskas, Andrea Setzer, Graham Katz, and Dragomir R. Radev. 2003. Timeml: Robust specification of event and temporal expressions in text. In *Fifth International Workshop on Computational Semantics*.

Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI*.

Hans Reichenbach. 1956. *The Direction of Time*. Dover Publications.

Nils Reimers and Iryna Gurevych. 2019. Sentence-bert: Sentence embeddings using siamese bert-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.

Ying Rex, Andrew Wang, Jiaxuan You, Chengtao Wen, Arquimedes Canedo, and Jure Leskovec. 2020. Neural subgraph matching. *arXiv*.

Mark Riedl and Robert Young. 2010. Narrative planning: Balancing plot and character. *Journal of Artificial Intelligence Research*, 39:217–268.

Andrew Rosenberg and Julia Hirschberg. 2007. V-measure: A conditional entropy-based external cluster evaluation measure. In *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Lang. Learning (EMNLP-CoNLL)*, pages 410–420, Prague. Association for Computational Linguistics.

Benedek Rozemberczki and Rik Sarkar. 2020. Characteristic Functions on Graphs: Birds of a Feather, from Statistical Descriptors to Parametric Models. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management (CIKM '20)*, page 1325–1334. ACM.

Keisuke Sakaguchi, Chandra Bhagavatula, Ronan Le Bras, Niket Tandon, Peter Clark, and Yejin Choi. 2021. proScript: Partially ordered scripts generation. In *Findings of the Assoc. for Computational Linguistics: EMNLP 2021*, pages 2138–2149, Punta Cana. Association for Computational Linguistics.

Roger C. Schank and Robert P. Abelson. 1977. *Scripts, plans, goals, and understanding : an inquiry into human knowledge structures*. Hillsdale, N.J. : Lawrence Erlbaum Associates.

Roger C. Schank and Robert P. Abelson. 1995. Knowledge and memory: The real story. In *Knowledge and Memory: the Real Story: Advances in Social Cognition, Volume VIII*, pages 1–85. Lawrence Erlbaum Associates.

Bernhard Schölkopf, Francesco Locatello, Stefan Bauer, Nan Rosemary Ke, Nal Kalchbrenner, Anirudh Goyal, and Yoshua Bengio. 2021. Toward causal representation learning. *Proceedings of the IEEE*, 109:612–634.

Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. Neural machine translation of rare words with subword units. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics*, pages 1715–1725, Berlin, Germany. Association for Computational Linguistics.

Leonard Talmy. 1988. Force dynamics in language and cognition. *Cognitive Science*, 12(1):49–100.

Fiona Anting Tan, Ali Hürriyetoğlu, Tommaso Caselli, Nelleke Oostdijk, Tadashi Nomoto, Hansi Hettiarachchi, Iqra Ameer, Onur Uca, Farhana Ferdousi Liza, and Tiancheng Hu. 2022. The causal news corpus: Annotating causal relations in event sentences from news. In *Proceedings of the Thirteenth Language Resources and Evaluation Conference*, pages 2298–2310, Marseille, France. European Language Resources Association.

Joshua B. Tenenbaum, Charles Kemp, Thomas L. Griffiths, and Noah D. Goodman. 2011. How to grow a mind: Statistics, structure, and abstraction. *Science*, 331(6022):1279–1285.

Jingxuan Tu, Kyeongmin Rim, Eben Holderness, and James Pustejovsky. 2022. Dense paraphrasing for textual enrichment. *CoRR*, abs/2210.11563.

Amos Tversky and Daniel Kahneman. 1982. Causal schemas in judgments under uncertainty. In Daniel Kahneman, Paul Slovic, and Amos Tversky, editors, *Judgment under Uncertainty: Heuristics and Biases*, page 117–128. Cambridge University Press.

Naushad UzZaman, Hector Llorens, Leon Derczynski, James Allen, Marc Verhagen, and James Pustejovsky. 2013. SemEval-2013 task 1: TempEval-3: Evaluating time expressions, events, and temporal relations. In *2nd Joint Conference on Lexical and Computational Semantics (*SEM)*, pages 1–9, Atlanta. Association for Computational Linguistics.

Sai Vallurupalli, Sayontan Ghosh, Katrin Erk, Niranjan Balasubramanian, and Francis Ferraro. 2022. POQue: Asking participant-specific outcome questions for a deeper understanding of complex events. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 8674–8697, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.

Petar Veličković, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph Attention Networks. *International Conference on Learning Representations*.

Xiaozhi Wang, Ziqi Wang, Xu Han, Wangyi Jiang, Rong Han, Zhiyuan Liu, Juanzi Li, Peng Li, Yankai Lin, and Jie Zhou. 2020. MAVEN: A Massive General Domain Event Detection Dataset. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 1652–1671, Online. Association for Computational Linguistics.

Haoyang Wen, Ying Lin, Tuan M Lai, Xiaoman Pan, Sha Li, Xudong Lin, Ben Zhou, Manling Li, Haoyu Wang, Hongming Zhang, Xiaodong Yu, Alexander Dong, Zhenhailong Wang, Yi R Fung, Piyush Mishra, Qing Lyu, Dídac Surís, Brian Chen, Susan W Brown, Martha Palmer, Chris Callison-burch, Carl Vondrick, Jiawei Han, Dan Roth, Shih-fu Chang, and Heng Ji. 2021. RESIN : A Dockerized Schema-Guided Cross-document Cross-lingual Cross-media Information Extraction and Event Tracking System. *NAACL*.

Peter West, Chandra Bhagavatula, Jack Hessel, Jena Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck, and Yejin Choi. 2022. Symbolic knowledge distillation: from general language models to commonsense models. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4602–4625, Seattle. Association for Computational Linguistics.

Phillip Wolff. 2007. Representing causation. *Journal of Experimental Psychology: General*, 136(1):82–111.

Ben Zhou, Kyle Richardson, Qiang Ning, Tushar Khot, Ashish Sabharwal, and Dan Roth. 2021. Temporal reasoning on implicit events from distant supervision. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1361–1371, Online. Association for Computational Linguistics.

A Appendix

A.1 Limitations

Limited scales of causation. We rely on a limited number and diversity of viewpoints for scales of causation (eight annotators at undergraduate and graduate levels at U.S. institutions of higher learning; four female, four male; five native, three non-native English speakers).

Size of dataset. In Table 1, we see that TORQUESTRA is slightly larger than other existing human-curated temporal and event structure reasoning datasets. Still, the question remains, how large must the corpus be to be able to enable successful learning of salient discourse-level explicit and implicit causal relations? It might be significant.

More and more varied test data. A collection of 3k Wikipedia articles in a single language is a relatively small sample for testing. We note the document label class imbalance: about 33% concern historical military conflicts. Also, labeled event types are noisy (e.g., "The area was <u>hit</u> by a storm" is labeled VIOLENCE-ATTACK). Topic labels were likely originally assigned using similar methods as our own (tf-idf).

Schema indeterminacy. For better schema matching models, it is possibly although not necessarily imperative to define in clear terms what a "schema" is. A "better" schema model is likely hierarchical in terms of event semantics, although the compositional nature of events is still indeterminate, and thus also generic events (schemas). Different manners of defining schemas are possible, some of which we explore in this work: as single event types, as sequences of event types (e.g., in a causal chain), as unordered sets of event types, as causal graphs, as having predicate argument structure, etc.

Causal relations and other possible experiments. We note that we have not reported on experiments with the causal subrelations shown in Table 2, work we leave for the future. We also have not reported on experiments ranking causal chain salience, nor on generation results directly conditioning on S_{causal} (limiting the use of causal schema graphs to matching experiments), a promising research direction. Data for each of these proposed experiments is available as part of our release.

Role of structured representations in NLP. The role of structured representations in NLP is debatable in the deep learning era (the importance of dependency parsing, syntactic treebanks, semantic role labeling, etc.). One concern is that such human annotation is expensive. Another concern is that in many cases linguistic structure can be preempted, using raw text only for both inputs and outputs.

To add to this debate, the success of methods in causal inference is largely based on the use of graphical structures, though there has not been yet a direct link made between how to automatically acquire these structures from language text using NLP methods, a concern we hope the present work helps address.
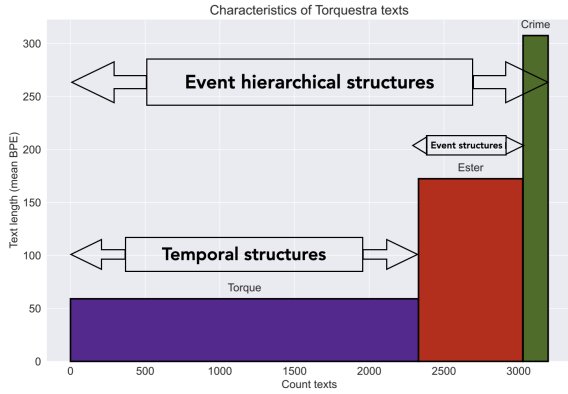
A.2 Describing Torquestra

Figure 6: Count of texts for different slices of TORQUESTRA by text mean length (#subword tokens). Existing human-engineered temporal and event structures (in boxes with arrows) supplement our causal structures (see Fig. 7 on right).
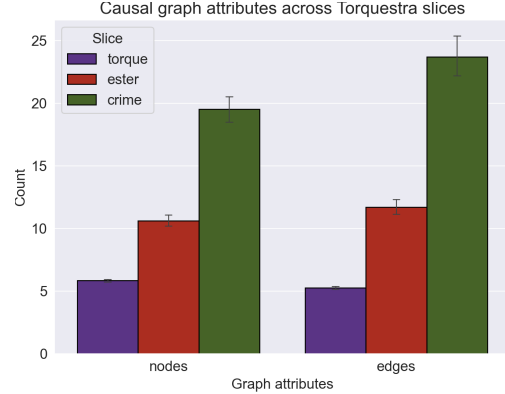

Figure 7: Complexity of human-engineered causal graphs in TORQUESTRA. The TORQUE data slice (purple, leftmost) consists of less complex causal graphs (based on number of nodes and edges) compared to CRIME (green, rightmost).

Text and graph statistics

Data attribute	$G_{temp\text{-}madaan}$	$G_{temp\text{-}torque}$	$G_{causal\text{-}torque}$	$G_{causal\text{-}ester}$	$G_{causal\text{-}crime}$
avg # tokens/text	120.3	46.8	46.8	133.4	231.3
max # nodes		22.0	17.0	19.0	36.0
mean # nodes	3.3	8.3	5.81	11.6	19.5
std # nodes		3.2	2.3	3.27	5.93
max # edges		102.0	21.0	24.0	48.0
mean # edges	3.9	20.9	5.2	13.1	23.7
std # edges		13.4	2.9	4.3	8.6
mean # enables			4.3	11.5	21.1
mean # blocks			0.9	1.6	2.6
mean degree			1.7	2.3	2.4
mean clustering			0.04	0.08	0.07
mean transitivity			0.05	0.09	0.07
mean sq clustering			0.02	0.05	0.03

Table 7: Comparison of data statistics from two temporal structure resources: $G_{temp\text{-}madaan}$ (Madaan and Yang, 2021) and $G_{temp\text{-}torque}$ (Ning et al., 2020) with various slices of TORQUESTRA: $G_{causal\text{-}torque}$, $G_{causal\text{-}ester}$ and $G_{causal\text{-}crime}$. One important comparison here is between the number of edges between G_{temp} (mean 20.9) versus $G_{causal\text{-}torque}$ (mean 5.2), which we could use to argue that causal structures provide a cleaner starting point for reasoning than temporal structures.

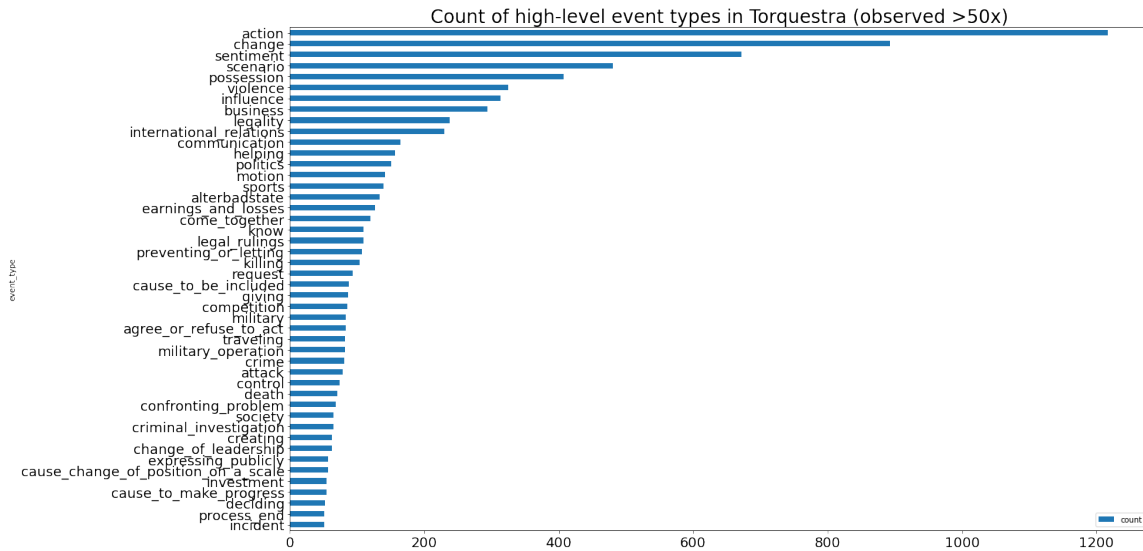## A.3 Analysis of events in TORQUESTRA

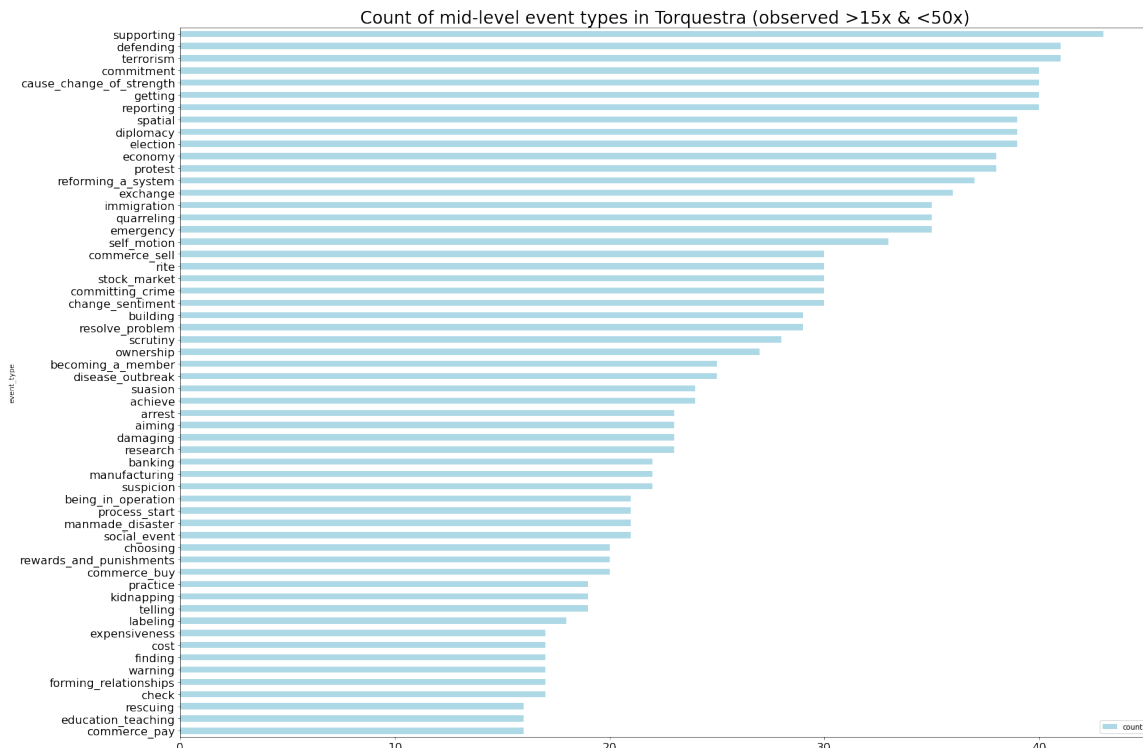Figure 8: Count of high-level hierarchical event semantic types observed in TORQUESTRA texts.

Figure 9: Count of mid-level hierarchical event semantic types observed in TORQUESTRA texts.

A.4 Annotation

General annotation guidelines. We refer to our approach of causal modeling as 'participant-centered causal structure' (§3). In this approach, nodes in causal graphs can be events expressed as sentences or phrases of Subject-Verb-Object form (e.g., 'Alice wrote a paper'). A range of other event descriptions are possible (regarding e.g., nominalizations, modality, polarity). Nodes can also be event participants such as people and objects that directly contribute to, initiate, or disrupt the beginning, unfolding, or ending of events. 'Alice made a cake' may be thus 'Alice' (as causal agent) → 'Alice made a cake'.

Annotators were asked to keep clearly in mind that one is not trying to model true causation, but just the apparent causation within the belief system of the speaker. Thus, causal links roughly mean "in the opinion of the speaker there should be a causal relation between A and B" — whether or not the actual causal relation (Table 2) is explicitly stated or implied somehow. This guideline helps remove the problem of reasoning about the world and allows one to frame debugging questions in terms of what the speaker might say and do, things that are much easier to discuss and evaluate.

Annotation quality. Consistent labeling of graph nodes proved challenging. For free-form node labeling, annotators were asked to compose a short sentence or phrase of form *Subject-Verb-Object*, e.g. 'Police arrested the suspect', or 'suspect arrested'. For evaluation, we selected 50 texts, asked two annotators to write short texts for nodes, counting the number that shared at least one content word. For the 325 nodes evaluated, 56% shared 1+ tokens referring to an event or other participant.

Notes about salient causal chain identification. Event salience detection is often framed as identifying the single most reportable event in an event sequence (Ouyang and McKeown, 2015). We extend this task to be: Given a set of triples, identify a subset of triples that tell the central part of the causal story. For evaluation, dual annotations for 50 causal graphs show agreement of $\kappa = 0.61$.

Notes about data evaluation: causal relations. For evaluation of human and machine generated causal graphs, annotators are given a text, two participants (**A** and **B**) in a directed relation, and three choices: A makes B more likely (ENABLES), less likely (BLOCKS) or has no effect, see Fig. 10. For 150 edges judged by three evaluators, average pairwise Cohen's kappa is $\kappa = 0.53$ (moderate agreement).

Figure 10: To evaluate a causal relation given a text and triple (rows), we used intuitive causal notions such 'A makes B more likely' in place of directly assessing an ENABLES edge.

A.5 Notes about data evaluation: causal graph completeness

We concede that causal graphs are difficult to fully specify for various reasons. First, causal expectations do not hold in certain contexts, implicit contributory factors related to events are typically beyond enumeration, and causal relations for events with negative polarity are easy to overlook.

Causal expectations for events that fail in a given context are difficult to agree on, e.g., subrelations WITHOUT EFFECT and UNKNOWN, ≈15% of relations in TORQUESTRA. In more than half of observed cases, explicit lexical cues (e.g., "despite" in "The protest happened despite the rain.") are not present.

Implicit contributory factors are numerous in certain kinds of text. In our case, newswire is necessarily succinct leaving a lot unsaid. Commonsense often suggests causal nodes of the kind "and etc.," perhaps shorthand for "there are many other potential contributory factors that I cannot all list, but they fit here." We leave the collection of this type of data for future work.

Notes about event polarity. Negative event polarity is a challenge for representation. There may be a reporting bias, i.e. perhaps people less frequently speak about the causal effects of negative polarity events, e.g., 'It didn't rain, so I didn't play in the mud'. Some of these semantics are captured in the WITHOUT EFFECTS causal subrelation. Annotators were instructed to look out for causal relations not covered by our ontology and to construct nodes including negative polarity, when possible. Additional thought and corpus studies may point to additional layers of meaning distinctions possible.

A.6 Prompt engineering

We construct various prompts for conditional language model generation. Using the pre-existing temporal and event structure annotations, prompts can be made concatenating raw texts with associated questions and answers about temporal and event structures. We also experiment by adding knowledge of hierarchical event types directly into texts, using dense paraphrasing (Tu et al., 2022).

Example prompt: text + G_{temp}. For prompts for text + G_{temp}, we randomly selected 2-3 temporal TORQUE questions and answers (Ning et al., 2020) in the form of lists of event mentions, filtering out verbs with less eventive meaning (e.g., 'was') and appended to the source text.

```
TEXT: And on that basis Keating was released from prison before he was eligible for parole.
Now the ninth US circuit court of appeals has ruled that the original appeal was flawed
since it brought up issues that had not been raised before.\n
What event has already finished? released, ruled, brought, flawed\n
What happened before the court ruled? flawed, brought, raised, released\n
What did not happen before Keating was released? parole\n
[GEN]
```

Example prompt: text + V_{event} (cf. dense paraphrasing). We also add structure to input text prepending event types to event and entity mentions. We wonder: Does it *help* or *hinder* generation by: 1) indicating which tokens are likely associated with causal chains? and, 2) giving an explicit signal of hierarchical event semantics? We also include event types for entities and events NOT part of causal chains (to aid in learning). Results are split on this question: in some cases, it helps, in other cases, it hinders (§6).

```
TEXT: And on that basis Entity::Keating was Legal_rulings::released from Entity::prison
before he was Scenario::eligible for parole.
Now the Entity::ninth US circuit court of appeals has Legal_rulings::ruled that
the original Legal_rulings::appeal was Incident::flawed
since it brought up Scenario::issues that had not been raised before.\n
[GEN]
```

A.7 Model specifics

We implemented graph generation models in PyTorch using the huggingface library. Without a limited hyperparameter search, we note generation models worked well with the following: block size between 300-600, training 12-16 epochs, learning rate=$3e^{-5}$, cosine learning rate scheduler, and Adam optimizer. For generation, we use nucleus sampling (Holtzman et al., 2019) setting $p = 0.90$ and temperature= 0.95.

To embed causal graphs, we implement a self-supervised Graph Attention Network (GAT) (Veličković et al., 2018) using PyTorch Geometric[8]. Our base model uses two GAT layers with ELU activations, dropout=0.5, global mean pooling and a final linear layer. We trained for 2-3 epochs over the entire dataset (TORQUESTRA + test data). Using the trained model, we make final embeddings for clustering experiments (test set alone); while for matching we compare the test set with TORQUESTRA graphs and RESIN subgraphs.

Masked objective function. For the self-supervised graph modeling, we train a graph neural network using message passing to compose graph embeddings, randomly masking input nodes. In training, we hypothesize that more important nodes will contribute more to a final graph embedding, and experiment masking nodes based on PageRank weighting[9].

Loss function. At each step of training, the model 'sees' 5 of 10 randomly sampled masked graphs, and predicts a graph embedding. We compare the similarity of the predicted embedding for the masked graph, $y*$, with the embedding for the entire ground truth graph, y and train the model to minimize an absolute value cosine similarity difference loss:

$$\mathcal{L}_{CosineSimDiff} = |1 - CosineSimilarity(y*, y)| \tag{2}$$

Cosine similarity, the scaled dot product of non-zero vectors **A** and **B**, is:

$$CosineSimilarity(\mathbf{A}, \mathbf{B}) = \frac{\mathbf{A} \cdot \mathbf{B}}{max(\|\mathbf{A}\|_2 \|\mathbf{B}\|_2, \epsilon)} \tag{3}$$

---

[8] https://pytorch-geometric.readthedocs.io/en/latest/
[9] Using the NetworkX Python package, https://networkx.org/

with $\epsilon = 1e^{-8}$ to prevent division by zero. For vectors that are equal, the cosine similarity is 1, 0 when orthogonal, and -1 when opposite, so the absolute value cosine similarity difference loss (Eq. 2) approaches zero as the embedding similarity for masked graph and non-masked graph increases. In our experiments, we mostly observed loss to decrease from 1 to about 0.6 in training and validation.

A.8 Metrics

For our experiments, we report a mixture of quantitative and qualitative metrics based on 2-3 human judgments. However, ground truth labels for nodes and schemas present a number of tricky issues (see Appendix A.1). This is apparent, for example, in judgments of generated graph correctness and also in the evaluation of schema matching models. Defining schemas as a single label is likely too simple, a situation that using a set of labels does not necessarily improve.

For schema matching, we report mean average precision (MAP), measured between 0 and 1 with higher being better. For Q queries, this is defined as:

$$MAP = \frac{1}{Q} \sum_{q=1}^{Q} \text{avg. precision}(q) \tag{4}$$

with precision for a query q defined as:

$$\text{avg. precision}(q) = \frac{|\{\text{relevant documents}\}_q \cap \{\text{retrieved documents}\}_q|}{|\{\text{retrieved documents}\}_q|} \tag{5}$$

We consider different quantitative and qualitative judgments for relevance, including: precision of topic label, lexical overlap and structural similarity, where structural similarity can be defined in terms of shared nodes, node degree, n-nary properties, and other observable graph properties. See a cluster example sharing structural similarity in Appendix B.

A.9 Data sample

Training examples are shorter texts from TORQUE (with temporal questions); Validation are longer texts from ESTER (with event structure questions) that **includes** the shorter TORQUE text, in **bold**.

```
{'split': 'train',
'source': 'torque',
'@id': 'train-docid_PRI19980115.2000.0186_sentid_6',
'notes': 'original-698',
'text': 'And on that basis Keating was released from prison before he was eligible for parole. Now the ninth US circuit court of appeals has ruled
    that the original appeal was flawed since it brought up issues that had not been raised before.',
'questions':
    ['What event has already finished?',
    'What event has begun but has not finished?',
    'What will happen in the future?',
    'What happened after Keating was released?',
    'What did not happen before Keating was released?',
    'What happened before the court ruled?',
    'What did not happen after Keating was released?',
    'What happened after the court ruled?'],
'answers':
    [['released', 'ruled', 'brought', 'flawed'], [], [],
    ['ruled', 'flawed', 'brought', 'raised'],
    ['was', 'parole'],
    ['flawed', 'brought', 'raised', 'released'], [], []],
'event_types':
    {'Keating was released': 'Action;Legality;Legal_rulings;Releasing',
    'Keating was eligible for parole': 'Action;Legality;Legal_rulings',
    'issues that had not been raised before': 'Scenario',
    'court ruled that the original appeal was flawed': 'Action;Legality;Legal_rulings'},
'noncausal_event_types':
    {'brought': 'Action;Communication;Reporting'},
'causal_graph': [
    {'head': 'court ruled that the original appeal was flawed', 'rel': 'ENABLES',
    'tail': 'Keating was released'},
    {'head': 'Keating was eligible for parole', 'rel': 'BLOCKS',
    'tail': 'Keating was released'},
    {'head': 'issues that had not been raised before', 'rel': 'ENABLES',
    'tail': 'ruled that the original appeal was flawed'}]}

{'split': 'dev',
'source': 'ester',
'@id': 'dev-docid_PRI19980115.2000.0186_sentid_6',
'notes': 'original-698',
'text': "Former savings and loan chief, Charles Keating, is facing more legal troubles in California.
    A federal appeals court has reinstated his state convictions for securities fraud.
    NPR's Elaine Corey has more from San Francisco. In nineteen ninety-one Charles Keating was convicted
    in state court of helping to defraud thousands of investors who bought high risk junk bonds
    sold by Keating's employees at Lincoln savings and loan. The bonds became worthless when the bankrupt thrift
    was seized by government regulators. Keating's convictions were thrown out in nineteen ninety-six on a technicality.
    And on that basis Keating was released from prison before he was eligible for parole. Now the ninth US circuit court of appeals has ruled
    that the original appeal was flawed since it brought up issues that had not been raised before.
    That means the convictions stand, a ruling likely to send Keating's lawyers
    back to state court where they must start over with a new appeal.",
'questions':
    ['Why was Mr. Keating convicted?', 'Why was Mr. Keating released from prison?',
    'What might happen as a result of the convictions being ruled to stand after the flawed appeal?'],
'answers':
    ['defraud thousands of investors',
    'convictions were thrown out in nineteen ninety-six on a technicality',
    "send Keating's lawyers back to state court"],
'event_types': 'banking;crime;action;legality;legal_rulings',
'causal_graph':
    [{'head': 'Entity::Charles Keating', 'rel': 'ENABLES',
    'tail': 'Keating faces legal troubles', 'saliency': 0},
    {'head': 'Entity::Charles Keating', 'rel': 'ENABLES',
    'tail': 'security fraud', 'saliency': 0},
    {'head': 'security fraud', 'rel': 'ENABLES',
    'tail': 'Keating is convicted of security fraud', 'saliency': 1},
    {'head': "Keating's convictions dismissed", 'rel': 'BLOCKS',
    'tail': 'Keating is convicted of security fraud', 'saliency': 1},
    {'head': "Keating's convictions dismissed", 'rel': 'ENABLES',
    'tail': 'Keating was released from prison', 'saliency': 1},
    {'head': 'technicality', 'rel': 'ENABLES',
    'tail': "Keating's convictions dismissed", 'saliency': 0},
    {'head': 'bonds became worthless', 'rel': 'ENABLES',
    'tail': 'security fraud', 'saliency': 0},
    {'head': 'bankrupt thrift was seized', 'rel': 'ENABLES',
    'tail': 'bonds became worthless', 'saliency': 0},
    {'head': 'Entity:ninth US circuit court of appeals', 'rel': 'ENABLES',
    'tail': "court reinstated Keating's state convictions for securities fraud", 'saliency': 1},
    {'head': "court reinstated Keating's state convictions for securities fraud", 'rel': 'ENABLES',
    'tail': 'Keating faces legal troubles', 'saliency': 1},
    {'head': "Keating's convictions dismissed", 'rel': 'ENABLES',
    'tail': "court reinstated Keating's state convictions for securities fraud", 'saliency': 1}]}
```

B Cluster example

Prototypical graph
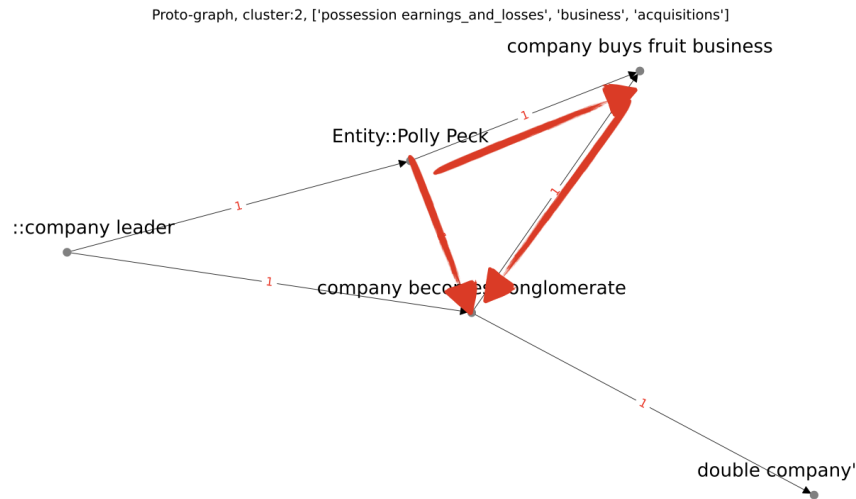

Figure 11: Example prototypical graph (cluster centroid) in an early experiment clustering TORQUESTRA*human* causal graphs.
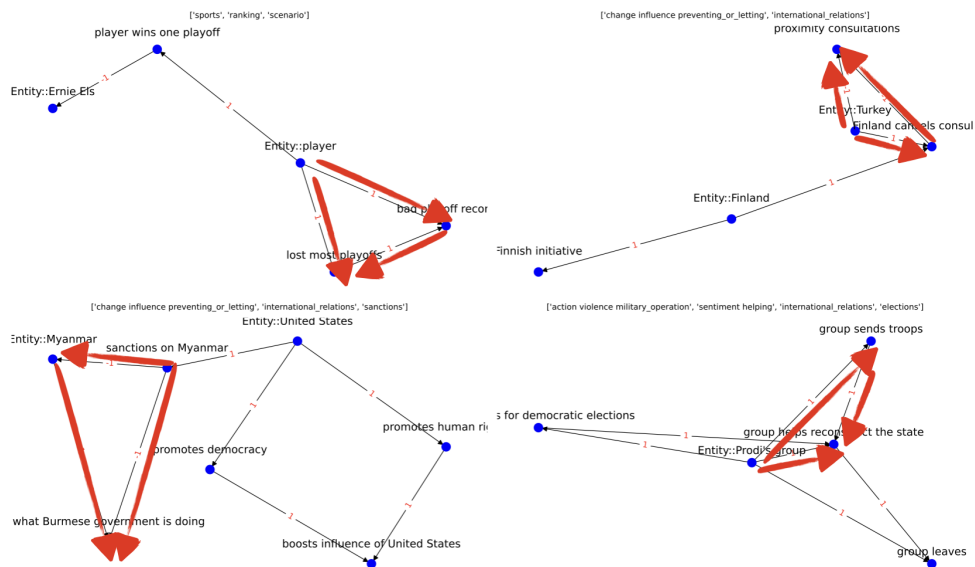
Top-4 similar graphs

Figure 12: Similar graphs from a cluster. The GNN matching algorithm is effective in identifying similar graph structures as well as those that share lexico-semantic similarity, cf. *motif discovery* (Milo et al., 2002).