

CoPERLex: Content Planning with Event-based Representations for Legal Case Summarization

Santosh T.Y.S.S, Youssef Farag, Matthias Grabmair

School of Computation, Information, and Technology;
Technical University of Munich, Germany

Abstract

Legal professionals often struggle with lengthy judgments and require efficient summarization for quick comprehension. To address this challenge, we investigate the need for structured planning in legal case summarization, particularly through event-centric representations that reflect the narrative nature of legal case documents. We propose our framework, CoPERLex, which operates in three stages: first, it performs content selection to identify crucial information from the judgment; second, the selected content is utilized to generate intermediate plans through event-centric representations modeled as Subject-Verb-Object tuples; and finally, it generates coherent summaries based on both the content and the structured plan. Our experiments on four legal summarization datasets demonstrate the effectiveness of integrating content selection and planning components, highlighting the advantages of event-centric plans over traditional entity-centric approaches in the context of legal judgements.

1 Introduction

Legal professionals, including lawyers, judges, and researchers, often face the daunting task of navigating lengthy legal judgments that are critical for case law interpretation and reasoning. Automatic legal summarization can significantly aid this process by producing concise overviews of extensive documents (Hachey and Grover, 2006). Earlier efforts centered on extractive summarization, valued for its faithfulness to the original text (Farzindar, 2004; Grover et al., 2003a,b; Saravanan et al., 2006; Bhat-tacharya et al., 2019; Polsley et al., 2016; Liu and Chen, 2019). However, recent advancements have shifted towards abstractive approaches (Shukla et al., 2022; Elaraby et al., 2023; Elaraby and Litman, 2022; Moro and Ragazzi, 2022), driven by the limitations of extractive methods, such as incomplete discourse, unresolved coreferences, and lack of context. Despite these, abstractive methods still

face challenges in capturing the complex context of legal documents, often leading to hallucinations, highlighting the need for more accurate techniques (Deroy et al., 2023b, 2024).

While most current approaches to legal case summarization rely on single end-to-end models to produce succinct, fluent, coherent, and factually accurate summaries, we argue that this setup often fails to effectively manage the implicit steps involved in summary generation, such as identifying salient content, organizing it logically, and maintaining faithfulness to the source. Moreover, this setup makes it difficult to mitigate specific types of errors without inadvertently introducing trade-offs that compromise other aspects of summary quality. These issues are particularly pronounced in the legal domain, where documents are often lengthy, dense, and rich in complex arguments and reasoning that require careful representation.

To address these limitations, we propose CoPERLex (Content Planning with Event-Based Representations for Legal Case Summarization), a three-stage framework inspired by how humans generate summaries from source documents. CoPERLex explicitly models the process as follows: (1) Content Selection, which involves analyzing the source document to identify the most salient content for inclusion in the summary; (2) Content Planning, where the selected content is organized into a structured intermediate representation, ensuring coherence and logical flow in the final summary; and (3) Summary Realization, which converts the intermediate representation into fluent, cohesive text.

For the **content selection** step, we employ an extractive summarization system, MemSum (Gu et al., 2022), which iteratively selects sentences based on a multi-step episodic Markov Decision Process, ensuring that the most relevant and salient content is identified for summarization. In the subsequent **content planning** stage, the selected salient content is used to produce an intermedi-

ate summary representation, ensuring coherence and logical structure for generating final summary. While various intermediate representations, such as entity chains (Narayan et al., 2021; Adams et al., 2024; Liu and Chen, 2021; Huot et al., 2024) or question-answer pairs as structural blueprints (Narayan et al., 2023; Huot et al., 2023), have been explored for short-text corpora, their application to lengthy legal judgments remains largely unexplored. Legal judgments, however, are inherently complex narratives that not only present facts but also describe a sequence of events, actions, and the legal reasoning applied to specific circumstances. This event-centric structure makes events crucial for understanding and summarizing such texts. In light of this, we seek to leverage event representations as an intermediate representation to guide the legal summarization process. We adopt a simpler representation of events in the form of Subject-Verb-Object (SVO) tuples to capture key actions and events, beyond any jurisdiction- or case-specific taxonomy, as carried out in Yao et al. (2022); Shen et al. (2020); Li et al. (2020), for better generalizability. We extract event triples from the golden summaries to augment the dataset with target plans and train an encoder-decoder-based Unlimiformer model (Bertsch et al., 2023) to generate the intermediate representation based on the extracted salient content.

Finally, in the **summary realization** stage, the intermediate event-based plan, combined with the selected content, is fed into a Longformer Encoder-Decoder model (Beltagy et al., 2020) to generate summaries. Our experiments on four legal summarization datasets—MultiLexSumm-Long, MultiLexSumm-Short, CanLII, and Super-SCOTUS—demonstrate the effectiveness of CoP-ERLex. The integration of content selection and planning yields significant improvements in faithfulness and coherence, with event-centric plans outperforming entity-centric representations.

2 Related Work

Legal Case Summarization: Earlier works have predominantly utilized extractive approaches (Bhattacharya et al., 2019) which encompass both unsupervised methods (Bhattacharya et al., 2021; Palsley et al., 2016; Farzindar, 2004; Saravanan et al., 2006; Mandal et al., 2021) and supervised methods (Agarwal et al., 2022; Liu and Chen, 2019; Zhong et al., 2019; Xu et al., 2021b; Man-

dal et al., 2021). Unsupervised methods for legal summarization include domain-independent techniques, such as frequency-based statistics (e.g., Luhn (Luhn, 1958)), graph-based methods (e.g., LexRank (Erkan and Radev, 2004)), and matrix factorization methods (e.g., LSA (Gong and Liu, 2001)). Additionally, legal-domain-specific approaches have been proposed, such as LetSum (Farzindar, 2004; Saravanan et al., 2006), which utilizes cue phrases to assign rhetorical and semantic roles to sentences in the source document and form a summary by selecting sentences, either through TF-IDF or K-mixture models. Palsley et al. (2016) use various domain-related features to score the importance of sentences, while (Zhong et al., 2019) creates template-based summaries and employs Maximal Marginal Relevance to handle redundancy in selected sentences. Bhattacharya et al. (2021) models summarization as an Integer Linear Programming optimization problem, incorporating guidelines from legal experts into the optimization setup. Additionally, Mandal et al. (2021) utilizes document-specific catchphrases that provide both domain-specific and document-specific information to guide the summarization process. Supervised methods employ diverse strategies, ranging from the knowledge engineering of different domain-specific features (Liu and Chen, 2019) to joint multi-task learning with Rhetorical Role Labeling (Agarwal et al., 2022) and leveraging the argument structure within documents (Xu et al., 2021b). Recently, Deroy et al. (2023a) evaluated an ensemble of existing extractive-based supervised and unsupervised methods, employing simple voting-based techniques as well as ranking-based and graph centrality approaches.

With the limitations of extractive approaches in providing an overall coherent summary (Zhang et al., 2022), recent works explored abstractive summarization techniques for legal case summarization, especially considering the recent successes of the unsupervised pre-training and supervised fine-tuning paradigm (Shukla et al., 2022; Ray et al., 2020; Schraagen et al., 2022; Santosh et al., 2024b; Tyss et al., 2024). Shukla et al. (2022) applied various transformer-based pre-trained abstractive models, such as BART, Legal-LED, and Legal-Pegasus for legal case summarization. To address the challenge posed by long legal documents, Moro and Ragazzi (2022) proposed chunking input documents into semantically coherent segments, summarizing each chunk individually and then con-

catenating the results to avoid truncation. Recent efforts have also focused on enhancing the factuality of generated summaries (Feijo and Moreira, 2023) by leveraging an entailment module to select faithful candidates. Additionally, rich argumentative structure embedded in legal documents are also explored to improve the performance of legal summarization models (Xu and Ashley, 2022; Elaraby and Litman, 2022; Elaraby et al., 2023). While most of these works concentrate on training and evaluating summarization models within specific jurisdictions, recent research by Santosh et al. (2024a) assesses the cross-jurisdiction generalizability of legal case summarization systems and further propose an adversarial learning approach to enhance cross-domain transfer of these methods.

To the best of our knowledge, no investigation has been conducted on planning-based approaches for legal summarization. In this work, we propose an event-centric representation as an intermediate plan for legal summarization.

Content Planning for Text Generation: Structured planning is considered a critical link in organizing content effectively before realization (Reiter and Dale, 1997), much like humans plan at a higher level than individual words, as evidenced by psycholinguistic studies (Levelt, 1993; Guhe, 2020). Earlier approaches incorporated various planning representations, such as Rhetorical Structure Theory (Mann and Thompson, 1988; Hovy, 1993) and MUC-style representations (McKeown and Radev, 1995), discourse trees (Mellish et al., 1998), entity transitions (Kibble and Power, 2004; Barzilay and Lapata, 2008), sequences of propositions (Karamanis, 2004), schemas (McKeown, 1985) and lexical chains (Barzilay and Elhadad, 1997).

Recent works in the data-to-text generation task divide it into two distinct phases: planning and realization of natural language text (Moryossef et al., 2019). Perez-Beltrachini and Lapata (2018) introduced content selection methods that do not rely on explicit planning. Meanwhile, Puduppully et al. (2019) and Laha et al. (2020) proposed micro-planning strategies, where they first establish a content plan based on a sequence of records and then generate a summary conditioned on that plan. Similar planning approaches have been explored for entity realization (Puduppully et al., 2019). Additionally, Puduppully and Lapata (2021) advocated for macro-planning as a means of organizing high-level document content, either in text form or

within latent space (Puduppully et al., 2022).

In summarization, various plan representations have been investigated. Narayan et al. (2020) treat step-by-step content selection as a planning component, generating sentence-level plans through extract-then-abstract methods (Zhang et al., 2022), and even sub-sentence-level plans using elementary discourse units (Adams et al., 2023). Further research by Narayan et al. (2021); Liu and Chen (2021) introduced intermediate plans using entity chains—ordered sequences of entities mentioned in the summary. More recently, Narayan et al. (2023); Huot et al. (2023) conceptualized text plans as a sequence of question-answer pairs, which serve as blueprints for content selection (i.e., what to say) and planning (i.e., in what order) in summarization tasks. Despite these advancements, there remains no consensus on the most effective representation.

In this work, we explore the utility of content planning approaches for summarizing long legal documents, specifically by designing an event-based plan representation that aligns with the narrative structure inherent to these texts.

Event-Centric Representations: Understanding events is critical for various tasks, including open-domain question answering (Yang et al., 2003), intent prediction (Rashkin et al., 2018), timeline construction (Do et al., 2012), text summarization (Daumé III, 2009), and misinformation detection (Fung et al., 2021). Events are not isolated predicates; they often interconnect through complex relationships, such as memberships, co-reference, temporal orders, and causal links. These relationships form narrative scripts, necessitating sophisticated representations that extend beyond simple predicates (Chen et al., 2021).

In the field of legal NLP, event-based representations have been underexplored for downstream tasks. However, recent work by Joshi et al. (2023) has utilized these representations for prior case retrieval, showcasing their potential. In this study, we leverage event-based representations to create an intermediate content plan specifically for legal case summarization. We adopt a simplified event representation from previous works (Chambers and Jurafsky, 2008, 2009; Jans et al., 2012), where events are represented as verb-dependency pairs (e.g., subject, object), in contrast to jurisdiction- or case-specific taxonomy, as carried out in Yao et al. (2022); Shen et al. (2020); Li et al. (2020), for better generalizability.

3 Our Method: CoPERLex

In this section, we present our proposed framework, CoPERLex, designed for the task of legal case summarization. Legal summarization requires processing complex and lengthy documents while ensuring that the generated summary is both coherent and faithful to the source. To address these challenges, we propose a pipelined approach that breaks down the task into three sub-tasks: (a) Content Selection Module identifies and selects the most salient sentences from the entire input document, similar to extractive summarization, ensuring that critical content is retained. (b) Content Planning Module takes the selected content as input and generates an intermediate content plan, serving as a blueprint for generating the final summary. This plan is constructed as event-based representations, expressed through subject-verb-object (SVO) triples. (c) Summary Generation Module synthesizes the summary by taking the selected content along with the event-based content plan as input, ensuring that the summary is structured and faithful to the outlined plan.

3.1 Content Selection

This module takes the input document segmented into sentences and outputs a set of salient sentences. We utilize the MemSum (Gu et al., 2022) (Multi-step Episodic Markov decision process extractive SUMmarizer), a reinforcement-learning-based extractive summarizer that iteratively selects sentences for the summary. At each time step, the policy network (agent) based on the state decides either to stop or continue to extract next sentence from remaining list of unselected sentences. If the agent continues, it computes a score for each remaining sentence and samples a sentence based on these scores. The state is represented by (a) the local content of the sentence (b) the global context of the sentence within the document and (c) the current extraction history (set of selected sentences till that point). To encode these three encodings in the state, we use a local sentence encoder, a global context encoder and an extraction history encoder. For local sentence encoder, we pass the word embeddings of each sentence through a bi-directional LSTM and multi-head pooling layer to obtain the sentence representation. The global context encoder takes each sentence embedding from local encoder and passes them through a bi-LSTM to obtain context-based embedding of each

sentence. Then extraction history encoder encodes the extraction history information and produces the extraction history embedding for each of the remaining unselected sentences through identical blocks, each containing two multi-head attention sublayers. The first sub-layer performs multi-head self-attention among the local embeddings of remaining sentences so that each remaining sentence can capture the context provided by other remaining sentences. The other attention sublayer performs multi-head attention over the embeddings of extracted sentences to enable each remaining sentence to attend to all the extracted sentences. This block finally outputs one history based embedding for each remaining sentence. Finally the local, global and extraction history based embedding of each unselected sentence is concatenated to produce score for each sentence. We also multi-head pool all these embeddings to obtain the stopping probability. This step-wise state-updating strategy allows the agent to consider the content of the partial summary during sentence selection, contrasting with single-step episodic MDP-based models (Dong et al., 2018; Narayan et al., 2018), which encode state information only once and extract the entire summary through a single action.

The policy network is trained using the Reinforce policy gradient algorithm (Williams, 1992), with the objective of maximizing the expected return at each time step. This return is defined as the cumulative reward from time $t + 1$ until the end of the episode, when the summary is complete. The instantaneous reward is zero, except at the end of the episode, where the final reward is computed as the mean of the ROUGE-1, 2, and L F1-scores between the extracted set of sentences (the summary) and the target abstractive summary. Notably, while training an extractive summarizer, we do not require explicit binary labels for each sentence as a supervision signal and instead utilize the target abstractive summary to calculate the reward.

3.2 Content Planning

We define the content plan as an ordered sequence of event-based representations. Each event is represented as a tuple containing a predicate (typically a verb or verb compound that describes the main action) and its primary arguments (e.g., subject, object, indirect object, and prepositional object). These tuples are typically structured as subject-verb-object triples. The content plan generator module utilizes the salient sentences from the previ-

ous module to construct a sequence of event tuples.

To train this module, both salient content and content plans for each document are required as input and output, respectively. However, since summarization corpora consist only of complete documents and their corresponding abstractive summaries, we need to extract sentence-level binary labels, which can be concatenated to form the salient content used as input for the plan generator module. Following [Liu and Chen \(2019\)](#), we select sentences for the summary based on their ROUGE-2 similarity score with the entire gold summary. This subset of sentences serves as the salient content for training our plan generator. During training, we utilize oracle-extracted salient sentences as input. In contrast, during inference, we rely on model-extracted salient sentences, which may be of lower quality and potentially introduce exposure bias in the plan generator. To mitigate this issue, we implement a hybrid training setup that combines outputs from the previous module with the golden extracted sentences, which enhances the model’s robustness against noisy extractions during inference.

To obtain the content plan from golden summary, we utilize Plumber ([Jaradeh et al., 2021](#)), a tool designed for creating knowledge graph triples from unstructured text. Plumber dynamically generates suitable information extraction pipelines, offering a total of 264 distinct pipelines. It employs a RoBERTa-based classification model to extract contextual embeddings from the input and identify an appropriate pipeline, which includes components for coreference resolution, triple extraction, entity linking, and relation linking to knowledge graphs. For our purposes, we focus primarily on coreference resolution and triplet extraction to derive the event representations for each sentence in the summary. These event representations are then concatenated in the order of the summary sentences to form the golden content plan.

Our plan generator employs an Unlimiformer model ([Bertsch et al., 2024](#)), a retrieval-based approach to augment pretrained language models to accept inputs of unbounded length at test time, to deal with longer inputs. Unlimiformer uses BART-encoder ([Lewis et al., 2020](#)) to encode overlapping chunks of the input and keep only the middle half of the encoded vectors from each chunk, to ensure that the encodings have sufficient context on both sides. Then it constructs a k-nearest-neighbor index over the hidden states of all input tokens. Then, every standard cross-attention head in every decoder

	Doc		Sum		Train	Valid	Test
	Tok.	Sent.	Tok.	Sen.			
MLS-Long	103712	3150	726	24	3177	454	908
MLS-Short	137042	4170	143	5	2210	312	616
CanLII	4914	174	291	11	839	104	106
SuperSCOTUS	5406	165	902	29	3246	406	406

Table 1: Statistics of datasets used in this paper.

layer queries the kNN index, such that the kNN distances are the attention dot-product scores, and attends only to the top-k input tokens. We train this model with our curated input-output sequences.

3.3 Summary Generation

This module takes the salient content selected by the content selection module and the generated event-based plan from the planning module to produce the final summary, ensuring coherence and adherence to the event-based narrative structure. We utilize oracle-extracted salient content and content plans derived from the target summary, as described in previous modules, to construct the input for training our summary generator. To address the exposure bias problem, we implement hybrid training, incorporating some inputs populated by model-extracted content and model-generated content plans, which enhances the model’s robustness during inference. Our summary generator employs a Longformer Encoder-Decoder (LED) ([Beltagy et al., 2020](#)) model, which employs local-global attention mechanism in the encoder to effectively handle longer input lengths.

4 Experiments

Datasets We experiment with the following four legal case summarization datasets:

MultiLexSum-Long/Short ([Shen et al., 2022](#)) comprises expert-written summaries for U.S. federal civil rights lawsuit cases, derived from multiple source documents such as complaints, settlement agreements, opinions, and orders. It features two summary granularities: (a) Long (L) summaries that detail the case background, parties, and events, and (b) Short (S) summaries that provide a concise overview in a single paragraph.

CanLII ([Xu et al., 2021a](#)) dataset includes legal case decisions and human-prepared summaries sourced from the Canadian Legal Information Institute, covering a variety of legal claims and issues presented in Canadian courts

SuperSCOTUS ([Fang et al., 2023](#)) is a multi-source

		R-1	R-2	R-L	BS	AS	Coh	Flu
MLS-Long	LED	47.83	24.95	28.69	86.31	64.22	75.50	78.79
	PRIMERA	45.14	23.68	27.47	85.46	62.48	75.83	78.23
	LongT5	46.26	24.22	27.76	85.82	63.73	75.79	78.78
	SLED-T5	46.73	24.04	28.24	86.34	65.18	75.12	77.66
	Unlimiformer-T5	48.42	24.56	28.11	86.18	66.24	74.82	77.53
	CoPERLEX	50.24*	26.57*	29.24*	86.92*	67.51*	77.84*	80.21*
MLS-Short	LED	45.66	22.46	30.64	87.94	61.45	68.15	81.46
	PRIMERA	42.27	21.68	29.18	86.93	61.57	69.40	81.88
	LongT5	44.24	22.04	29.87	87.68	61.25	69.21	81.39
	SLED-T5	45.52	22.27	30.12	88.12	61.87	68.03	80.65
	Unlimiformer-T5	45.46	23.13	30.77	88.05	62.51	67.64	80.76
	CoPERLEX	46.18*	23.04	31.13*	88.69*	63.49*	69.27	81.79
CanLII	LED	47.99	24.47	43.66	86.30	63.13	64.70	86.25
	PRIMERA	47.65	24.19	44.04	86.77	64.72	64.24	86.18
	LongT5	48.08	24.76	44.48	86.26	64.62	64.78	85.78
	SLED-T5	48.12	25.05	44.62	86.81	65.31	65.06	86.09
	Unlimiformer-T5	48.62	25.62	44.58	87.66*	65.06	64.92	86.14
	CoPERLEX	49.19	25.88	45.44*	87.25	66.23*	65.42*	86.53
SuperSCOTUS	LED	51.14	24.82	28.11	82.12	56.70	63.50	67.60
	PRIMERA	48.65	24.17	26.69	82.88	56.25	63.77	67.12
	LongT5	49.28	24.88	27.37	83.34	56.39	64.16	67.30
	SLED-T5	50.36	24.61	28.24	82.91	56.18	63.12	66.18
	Unlimiformer-T5	50.21	24.49	27.79	83.18	57.12	63.29	66.49
	CoPERLEX	52.28*	25.78*	28.94*	83.52*	59.38*	66.21*	69.49*

Table 2: Performance comparison of CoPERLex across four datasets. BS, AS, Coh, and Flu represent BERTScore, AlignScore, Coherence and Fluency, respectively. Entries marked with * indicate statistically significant improvement over the second-best performance, using the Wilcoxon signed-rank test with a 95% confidence interval.

dataset of U.S. Supreme Court cases, integrating resources such as oral arguments, post-hearing annotations, and summaries, including case Opinions. For our work, we utilize the summarization dataset, which generates the Syllabus from the Court’s majority opinions excluding concurring and dissenting opinions. The Syllabus, prepared by the Reporter of Decisions—a statutory officer under the Court’s direction—summarizes the case background, details from lower courts, and the court’s reasoning.

Detailed statistics for these datasets are presented in Table 1. The MLS-Long and Short datasets feature longer input lengths, highlighting their multi-document nature. Both MLS-Long and SuperSCOTUS exhibit longer summary lengths, while MLS-Short presents a greater challenge due to its higher compression ratio.

Metrics: We evaluate the generated summaries using ROUGE-1, 2, and L (Lin, 2004) to assess

lexical overlap with reference paragraphs, and BERTScore (Zhang et al., 2019a) for semantic similarity between the generated and reference summaries. For faithfulness, we use AlignScore (Zha et al., 2023) to measure factual consistency through a unified alignment function between the input context and generated text. Additionally, coherence and fluency scores are reported using the UniEval metric (Zhong et al., 2022).

Baselines: We compare our CoPERLex approach with several abstractive summarization models designed for long documents: (i) LED (Beltagy et al., 2020) (Longformer Encoder-Decoder), (ii) PRIMERA (Xiao et al., 2022), (iii) LongT5 (Guo et al., 2022), (iv) SLED (Ivgi et al., 2023), and (v) Unlimiformer (Bertsch et al., 2023). Detailed descriptions of the baselines and implementation specifics for CoPERLex can be found in App. A and B respectively.

	MLS-Long							CanLII						
	R-1	R-2	R-L	BS	AS	Coh	Flu	R-1	R-2	R-L	BS	AS	Coh	Flu
CoPERLex	50.24	26.57	29.24	86.92	67.51	77.84	80.21	49.19	25.88	45.44	87.25	66.23	65.42	86.53
w/o Content Sel.	48.14	25.14	28.54	86.42	66.03	75.34	79.78	49.06	25.41	45.28	86.72	66.14	65.72	86.18
w/o Planning	49.98	25.22	28.91	86.44	66.64	75.84	79.27	48.17	24.22	44.09	86.37	64.36	65.05	86.27
w/o CS & Planning	47.83	24.95	28.69	86.31	64.22	75.50	78.79	47.99	24.47	43.66	86.30	63.13	64.70	86.25
w/o Hybrid Train.	48.61	25.36	28.66	86.45	63.87	76.12	79.16	48.38	25.03	44.81	86.54	63.29	65.12	86.18

Table 3: Ablation study to assess effect of different components in CoPERLex framework. Con. Sel. (CS) & Pl. denote Content Selection & Planning components respectively.

4.1 Results

From Table 2, we observe that LED consistently outperforms PRIMERA across all datasets, which can be attributed to the larger input length of LED (16k vs. 4k). This highlights the absence of lead bias in legal datasets, unlike news datasets, and underscores the need for better long-context modeling strategies. LongT5 also surpasses PRIMERA, likely due to its extended input length, but still falls short of LED’s performance.

SLED performs comparably to these long-range models on MLS-Long, MLS-Short, and SuperSCOTUS, while outperforming them on CanLII. Unlimiformer generally outperforms all the long-range models except on SuperSCOTUS, demonstrating that using off-the-shelf, short-range pre-trained models like BART within the SLED or Unlimiformer framework yields competitive results without the need for long-range pre-training or specialized summarization objectives. Additionally, Unlimiformer’s strategy of attending to the top-k input keys proves to be an effective approximation of full attention. However, while SLED and Unlimiformer show competitive performance in content-based metrics like ROUGE, BERTScore, and AlignScore, they underperform in coherence and fluency across most datasets, except for CanLII in terms of coherence. This suggests that approximate attention and short-range models affect stylistic elements, indicating the need for improved approaches to enhance these aspects.

Overall, our proposed CoPERLex consistently outperforms all tested models across all datasets in most of the content and style metrics. This can be attributed to its modular framework handling three complementary steps of content selection to filter irrelevant information, event-centric outline generation to structure the narrative and ensure coherence and final synthesis to combine the outline with selected content into a cohesive summary.

This division prevents the end-to-end model from becoming overwhelmed, allowing it to produce more concise, coherent, and contextually accurate outputs. The improvements are higher in MLS-Long and SuperSCOTUS, which feature longer summaries (Table 1), particularly evident in coherence, fluency, and AlignScore, as end-to-end models struggle to balance factual accuracy with structural integrity. CoPERLex, by managing sub-tasks independently, performs each step more effectively. For datasets with shorter summaries, such as CanLII and MLS-Short, end-to-end models are comparable or even superior in style-based metrics. However, CoPERLex excels in factual accuracy, as it better handles the lengthy nature of inputs with content selection, where end-to-end models often fail to preserve source fidelity, demonstrating the advantage of multi-step approach.

Ablation Study To assess the contribution of each component in CoPERLex, we conducted an ablation study by testing four variations: (a) without the content selection component (w/o Con. Sel.), where we use the entire document to produce a plan and generate the summary, bypassing content filtering; (b) without the planning component, where we directly extract salient content and generate the summary in a traditional extractive-then-abstractive approach; (c) without both the content selection and planning components, using the entire document to produce a summary without any intermediate steps and (d) without Hybrid Training in both the plan generation and summary generation modules. The results for these variants, evaluated on the MLS-Long and CanLII datasets, are presented in Table 3. Our findings reveal that removing both components leads to the most significant performance degradation. For MLS-Long, the removal of content selection had the greatest impact, as the longer input made it more challenging for the planning model to both understand salient content and

generate a coherent summary structure. In contrast, for CanLII, removing the planning component was more detrimental, as the shorter input meant the absence of content selection had less of an effect. The lack of a planning results in disjointed summaries that fail to reflect the logical progression of legal arguments or case facts, evidenced by coherence, align scores. Finally, removing the hybrid training mechanism degrades CoPERLex’s performance, emphasizing the robustness offered by this.

Plan Representation We evaluated various representations for the planning component: (a) entity-centric plans, represented by entity chains—ordered sequences of entities mentioned in the summary, as proposed by Narayan et al. (2021); (b) event-centric plans, where events are represented as tuples of subject, verb, and object, extracted through dependency parsing to capture syntactic relationships. This approach, used in prior work by Joshi et al. (2023) for legal case retrieval, is adapted here for summarization; (c) event-centric plans from Plumber (Jaradeh et al., 2021), a framework that uses a multi-pipeline approach to extract entities, relations, and handle coreferences, which we adopt in CoPERLex. We provided the summarization model with oracle plans (extracted from the reference summaries) along with the source documents, and reported the results on the MLS-Long dataset in Table 4. Our findings show that event-centric plans consistently outperform entity-centric ones, reflecting the unique narrative structure of legal documents. While a limited number of entities (such as the involved parties) remain central, the complexity of the events surrounding them intensifies over time. This growing complexity is essential for capturing the evolving arguments, decisions, and legal reasoning in a case. Event-centric representations are thus more effective for summarizing legal documents, as they focus on dynamic interactions among entities rather than static entities. Among event-centric methods, Plumber’s orchestration of pipelines for entity and relation extraction delivers superior representations, allowing for a more accurate and coherent depiction of legal discourse, resulting in improved performance.

Content Selection We analyze various extractive algorithms for the content selection component, comparing MemSum with BERT-Ext (Miller, 2019) and HiBERT (Zhang et al., 2019b). BERT-Ext is an unsupervised approach that embeds source documents and selects k key sentences from the embedding clusters, while HiBERT utilizes

	R-1	R-2	R-L	BS	AS	Coh
Entities	54.14	30.16	38.56	88.66	77.13	80.45
UCREAT	57.72	32.61	41.91	89.61	78.46	80.65
Plumber	59.65	36.16	44.26	89.72	80.71	82.68

Table 4: Comparison of various planning approaches on the MLS-Long dataset, using oracle plans derived from the reference summary as input to the summary generation module.

	MLS-Long			CanLII		
	R-1	R-2	R-L	R-1	R-2	R-L
BERT-Ext	34.63	9.10	13.17	38.43	23.75	20.56
HiBERT	36.28	11.28	14.83	41.27	24.74	26.75
MemSum	36.92	12.85	15.56	42.16	25.19	28.16

Table 5: Comparison of different extractive approaches for Content Selection.

a hierarchical document encoder pre-trained on longer contexts using unlabeled data. We report the performance of the extractive summary against the abstractive reference summary using ROUGE scores in Table 5 on MLS-Long dataset. The superior performance of MemSum can be attributed to two main factors: (1) its stopping criterion is modeled as an action in RL framework that reduces the necessity to predefine the number of required sentences, and (2) its effective use of extraction history, allowing for iterative sentence selection rather than relying on a one-step extraction process.

Plan Generation In the plan generation phase, we utilize the entire document as input for the planning model to produce event-based tuples as outlines. We evaluate the performance of the LED model and the Unlimiformer model by measuring ROUGE scores between the generated plans and the extracted plans from the reference summary. As shown in Table 6, Unlimiformer consistently outperforms LED, particularly for longer sequences in R-2 and R-L. Consequently, we chose Unlimiformer as our plan generator in CoPERLex.

	Plan Gen.			Sum. Gen. with OP		
	R-1	R-2	R-L	R-1	R-2	R-L
LED	35.18	15.03	20.81	59.65	36.16	44.26
Unlimiformer	34.96	15.65	21.44	57.16	33.88	41.18

Table 6: Comparison of model architectures for Plan Generation and Summary Generation using Oracle Plans (OP) on MLS-Long dataset.

Summary Generation For the summary generation, we use oracle plans extracted from the reference summary as input to the summarizer model, alongside the document. Again, we explore the performance of both LED and Unlimiformer models. Table 6 indicates that LED significantly outperforms Unlimiformer in this context, which contrasts with our earlier findings in Table 2 on the Plan Generation section. We attribute this discrepancy to Unlimiformer’s kNN-based attention mechanism, which may struggle to adhere to the provided plan triples due to its lack of an exclusive focus mechanism. This limitation hinders its ability to concentrate on planning tuples while retrieving k tokens during decoding. Therefore, we opted to use LED as our summary generator in CoPERLex.

Case Study: In this section, we explore two key advantages of our event-based content planning framework: (i) how it enhances structural coherence and faithfulness to the source document and (ii) how it offers controllability for generating more customized summaries. These aspects are illustrated with examples in App. C.

Without a content plan, summaries often lack structure and coherence, leading to irrelevant or incomplete information. They may even fabricate details that are not aligned with the source document. As seen in Table 8, key legal actions, such as the filing of a lawsuit, are omitted, resulting in incomplete summaries that fail to capture critical events. The date of the filing is fabricated as the date of settling the lawsuit, creating unfaithful summaries. Entity-based planning offers some improvement by grounding the summary in entities or crucial dates, which reduces the likelihood of fabrication. However, as observed, focusing solely on entities can still miss essential actions and relationships, leading to unreliable summaries, such as misinterpreting the filing date as the intervention date. This occurs because events often revolve around entities, and without a clear grounding strategy, multiple possibilities arise for interpreting those entities, making it challenging to create a faithful summary. Our event-based planning approach addresses these challenges by grounding the summary in specific actions and interactions, rather than just the entities, making the summary more anchored to the source document.

In Table 7, we modified the event sequence by incorporating relevant details (such as the ‘requirement for signatures from 5% of registered voters’) and removing irrelevant elements (like ‘seeking re-

lief’). This adjustment affected the final summary, showcasing how users can customize content for improved conciseness and informativeness through event-based planning. Additionally, this process illustrates a potential method for enhancing summary faithfulness by employing similar strategies to develop automated methods for verifying and refining unfaithful event tuples, leading to more reliable outputs.

5 Conclusion

In this paper, we introduced CoPERLex, a novel approach for legal case summarization that integrates salient content selection and content planning with event-centric representations. Our experiments showed that utilizing structured event-based plans significantly enhances coherence and narrative flow, outperforming state-of-the-art models across four legal datasets. By capturing relationships between entities within legal texts, CoPERLex enhances summarization quality and paves the way for future research in legal NLP, underscoring the importance of event-driven frameworks in addressing the complexity of legal narratives.

Limitations

While CoPERLex employs event-centric representations to enhance narrative flow in legal case summarization, its current approach oversimplifies the representation of events and their interrelationships. Legal cases often involve intricate temporal relations, event co-references, and complex cause-and-effect chains, necessitating more sophisticated narrative modeling. Future work could improve CoPERLex by incorporating advanced representations that effectively capture these temporal dependencies, as well as utilizing graph-based structures to represent the interplay among multiple events and legal arguments. This enhancement would enable the system to better address more complex legal cases, where the sequence and interaction of events are pivotal to case outcomes.

Another limitation of this work lies in the evaluation methodology. While CoPERLex’s performance is assessed using widely established metrics such as ROUGE, BERTScore, and AlignScore, these metrics are limited in their ability to fully capture the unique complexities of legal texts, such as legal reasoning and the argumentative structure. Future research could focus on developing domain-specific evaluation metrics tailored to the needs of

legal summarization, potentially involving more qualitative measures of reasoning accuracy and argument structure. Additionally, the lack of validation by legal experts is a notable limitation of this study. Direct feedback from legal professionals could provide invaluable insights into the practical utility of the summaries generated by CoPERLex, but this was not feasible due to the unavailability of legal experts for evaluation.

Finally, CoPERLex has primarily been evaluated on English-language legal datasets, which limits its applicability to other languages. Extending the model to non-English legal documents presents an interesting avenue for future exploration. Moreover, its evaluation has been conducted in a same-jurisdiction setting—i.e., trained and tested on cases from the same jurisdiction. Given the diverse nature of legal systems across countries and the limited availability of annotated summarization datasets, ensuring cross-jurisdictional generalizability of summarization systems becomes essential (Santosh et al., 2024a). Investigating CoPERLex in cross-jurisdictional settings remains another interesting future direction.

Ethics Statement

The datasets utilized in this research were either publicly accessible or obtained through agreements with their respective owners, ensuring ethical compliance in their usage. While these datasets are not anonymized, they comprise legal documents that are already in the public domain or made available through official channels. As a result, their use in our study does not pose any additional harm beyond their existing availability.

The development of legal summarization systems like CoPERLex presents several ethical challenges. First, while automating the summarization of legal documents can significantly reduce the workload for legal professionals, it is crucial to recognize that such systems are not substitutes for human expertise, particularly in sensitive legal contexts. The summaries generated by CoPERLex may fail to capture all the nuances of complex legal cases, which could lead to misinterpretation or oversimplification of critical details. To address this concern, the model’s outputs should be regarded as assistive tools rather than authoritative sources and must be thoroughly reviewed by qualified legal professionals before any formal application.

Second, the biases inherent in the training data

may influence the fairness and objectivity of the generated summaries. Legal datasets often reflect societal biases and historical injustices, raising the risk that CoPERLex could inadvertently perpetuate these biases in its outputs. Therefore, it is essential to implement careful selection of training data, continuous monitoring, and effective bias mitigation strategies to minimize these risks.

References

- Griffin Adams, Alexander R Fabbri, Faisal Ladhak, Kathleen McKeown, and Noémie Elhadad. 2023. Generating edu extracts for plan-guided summary re-ranking. In *Proceedings of the conference. Association for Computational Linguistics. Meeting*, volume 2023, page 2680. NIH Public Access.
- Griffin Adams, Jason Zucker, and Noémie Elhadad. 2024. Speer: Sentence-level planning of long clinical summaries via embedded entity retrieval. *arXiv preprint arXiv:2401.02369*.
- Abhishek Agarwal, Shanshan Xu, and Matthias Grabmair. 2022. Extractive summarization of legal decisions using multi-task learning and maximal marginal relevance. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 1857–1872.
- Joshua Ainslie, Santiago Ontanon, Chris Alberti, Václav Cvicek, Zachary Fisher, Philip Pham, Anirudh Ravula, Sumit Sanghai, Qifan Wang, and Li Yang. 2020. Etc: Encoding long and structured inputs in transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 268–284.
- Regina Barzilay and Michael Elhadad. 1997. Using lexical chains for text summarization.
- Regina Barzilay and Mirella Lapata. 2008. Modeling local coherence: An entity-based approach. *Computational Linguistics*, 34(1):1–34.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Amanda Bertsch, Uri Alon, Graham Neubig, and Matthew Gormley. 2024. Unlimiformer: Long-range transformers with unlimited length input. *Advances in Neural Information Processing Systems*, 36.
- Amanda Bertsch, Uri Alon, Graham Neubig, and Matthew R Gormley. 2023. Unlimiformer: Long-range transformers with unlimited length input. *arXiv preprint arXiv:2305.01625*.
- Paheli Bhattacharya, Kaustubh Hiware, Subham Rajgaria, Nilay Pochhi, Kripabandhu Ghosh, and Saptarshi Ghosh. 2019. A comparative study of summarization algorithms applied to legal case judgments.

- In *Advances in Information Retrieval: 41st European Conference on IR Research, ECIR 2019, Cologne, Germany, April 14–18, 2019, Proceedings, Part I 41*, pages 413–428. Springer.
- Paheli Bhattacharya, Soham Poddar, Koustav Rudra, Kripabandhu Ghosh, and Saptarshi Ghosh. 2021. Incorporating domain knowledge for extractive summarization of legal case documents. In *Proceedings of the eighteenth international conference on artificial intelligence and law*, pages 22–31.
- Nathanael Chambers and Dan Jurafsky. 2008. Unsupervised learning of narrative event chains. In *Proceedings of ACL-08: HLT*, pages 789–797.
- Nathanael Chambers and Dan Jurafsky. 2009. Unsupervised learning of narrative schemas and their participants. In *Proceedings of the Joint Conference of the 47th Annual Meeting of the ACL and the 4th International Joint Conference on Natural Language Processing of the AFNLP*, pages 602–610.
- Muhao Chen, Hongming Zhang, Qiang Ning, Manling Li, Heng Ji, Kathleen McKeown, and Dan Roth. 2021. Event-centric natural language processing. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: Tutorial Abstracts*, pages 6–14.
- Hal Daumé III. 2009. Bayesian query-focused summarization. *arXiv preprint arXiv:0907.1814*.
- Aniket Deroy, Kripabandhu Ghosh, and Saptarshi Ghosh. 2023a. Ensemble methods for improving extractive summarization of legal case judgements. *Artificial Intelligence and Law*, pages 1–59.
- Aniket Deroy, Kripabandhu Ghosh, and Saptarshi Ghosh. 2023b. How ready are pre-trained abstractive models and llms for legal case judgement summarization? *arXiv preprint arXiv:2306.01248*.
- Aniket Deroy, Kripabandhu Ghosh, and Saptarshi Ghosh. 2024. Applicability of large language models and generative models for legal case judgement summarization. *Artificial Intelligence and Law*, pages 1–44.
- Quang Do, Wei Lu, and Dan Roth. 2012. Joint inference for event timeline construction. In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 677–687.
- Yue Dong, Yikang Shen, Eric Crawford, Herke van Hoof, and Jackie Chi Kit Cheung. 2018. Banditsum: Extractive summarization as a contextual bandit. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3739–3748.
- Mohamed Elaraby and Diane Litman. 2022. Arglegal-sum: Improving abstractive summarization of legal documents with argument mining. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6187–6194.
- Mohamed Elaraby, Yang Zhong, and Diane Litman. 2023. Towards argument-aware abstractive summarization of long legal opinions with summary reranking. *arXiv preprint arXiv:2306.00672*.
- Günes Erkan and Dragomir R Radev. 2004. Lexrank: Graph-based lexical centrality as salience in text summarization. *Journal of artificial intelligence research*, 22:457–479.
- Biaoyan Fang, Trevor Cohn, Timothy Baldwin, and Lea Frermann. 2023. Super-scotus: A multi-sourced dataset for the supreme court of the us. In *Proceedings of the Natural Legal Language Processing Workshop 2023*, pages 202–214.
- Atefeh Farzindar. 2004. Atefeh farzindar and guy lapalme; letsum, an automatic legal text summarizing system’ in t. gordon (ed.), legal knowledge and information systems. jurix 2004: The seventeenth annual conference. amsterdam: Ios press, 2004, pp. 11-18. In *Legal knowledge and information systems: JURIX 2004, the seventeenth annual conference*, volume 120, page 11. IOS Press.
- Diego de Vargas Feijo and Viviane P Moreira. 2023. Improving abstractive summarization of legal rulings through textual entailment. *Artificial intelligence and law*, 31(1):91–113.
- Yi Fung, Christopher Thomas, Revanth Gangi Reddy, Sandeep Polisetty, Heng Ji, Shih-Fu Chang, Kathleen McKeown, Mohit Bansal, and Avirup Sil. 2021. Infosurgeon: Cross-media fine-grained information consistency checking for fake news detection. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1683–1698.
- Yihong Gong and Xin Liu. 2001. Generic text summarization using relevance measure and latent semantic analysis. In *Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 19–25.
- Claire Grover, Ben Hachey, Ian Hughson, and Chris Korycinski. 2003a. Automatic summarisation of legal documents. In *Proceedings of the 9th international conference on Artificial intelligence and law*, pages 243–251.
- Claire Grover, Ben Hachey, and Chris Korycinski. 2003b. Summarising legal texts: Sentential tense and argumentative roles. In *Proceedings of the HLT-NAACL 03 Text Summarization Workshop*, pages 33–40.
- Nianlong Gu, Elliott Ash, and Richard Hahnloser. 2022. Memsum: Extractive summarization of long documents using multi-step episodic markov decision

- processes. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6507–6522.
- Markus Guhe. 2020. *Incremental conceptualization for language production*. Psychology Press.
- Mandy Guo, Joshua Ainslie, David C Uthus, Santiago Ontanon, Jianmo Ni, Yun-Hsuan Sung, and Yinfei Yang. 2022. Longt5: Efficient text-to-text transformer for long sequences. In *Findings of the Association for Computational Linguistics: NAACL 2022*, pages 724–736.
- Ben Hachey and Claire Grover. 2006. Extractive summarisation of legal texts. *Artificial Intelligence and Law*, 14:305–345.
- Eduard H Hovy. 1993. Automated discourse generation using discourse structure relations. *Artificial intelligence*, 63(1-2):341–385.
- Fantine Huot, Joshua Maynez, Chris Alberti, Reinald Kim Amplayo, Priyanka Agrawal, Constanza Fierro, Shashi Narayan, and Mirella Lapata. 2024. μ -plan: Summarizing using a content plan as a cross-lingual bridge. In *Proceedings of the 18th Conference of the European Chapter of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2146–2163.
- Fantine Huot, Joshua Maynez, Shashi Narayan, Reinald Kim Amplayo, Kuzman Ganchev, Annie Priyadarshini Louis, Anders Sandholm, Dipanjan Das, and Mirella Lapata. 2023. Text-blueprint: An interactive platform for plan-based conditional generation. In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics: System Demonstrations*, pages 105–116.
- Maor Ivgi, Uri Shaham, and Jonathan Berant. 2023. Efficient long-text understanding with short-text models. *Transactions of the Association for Computational Linguistics*, 11:284–299.
- Bram Jans, Steven Bethard, Ivan Vulic, and Marie-Francine Moens. 2012. Skip n-grams and ranking functions for predicting script events. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics (EACL 2012)*, pages 336–344. ACL; East Stroudsburg, PA.
- Mohamad Yaser Jaradeh, Kuldeep Singh, Markus Stocker, Andreas Both, and Sören Auer. 2021. Better call the plumber: orchestrating dynamic information extraction pipelines. In *International Conference on Web Engineering*, pages 240–254. Springer.
- Abhinav Joshi, Akshat Sharma, Sai Kiran Tanikella, and Ashutosh Modi. 2023. U-creat: Unsupervised case retrieval using events extraction. *arXiv preprint arXiv:2307.05260*.
- Nikiforos Karamanis. 2004. Entity coherence for descriptive text structuring.
- Rodger Kibble and Richard Power. 2004. Optimizing referential coherence in text generation. *Computational Linguistics*, 30(4):401–416.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Anirban Laha, Parag Jain, Abhijit Mishra, and Karthik Sankaranarayanan. 2020. Scalable micro-planned generation of discourse from structured data. *Computational Linguistics*, 45(4):737–763.
- Willem JM Levelt. 1993. *Speaking: From intention to articulation*. MIT press.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880.
- Qingquan Li, Qifan Zhang, Junjie Yao, and Yingjie Zhang. 2020. Event extraction for criminal legal text. In *2020 IEEE international conference on knowledge graph (ICKG)*, pages 573–580. IEEE.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Chao-Lin Liu and Kuan-Chun Chen. 2019. Extracting the gist of chinese judgments of the supreme court. In *proceedings of the seventeenth international conference on artificial intelligence and law*, pages 73–82.
- Zhengyuan Liu and Nancy F Chen. 2021. Controllable neural dialogue summarization with personal named entity planning. *arXiv preprint arXiv:2109.13070*.
- Hans Peter Luhn. 1958. The automatic creation of literature abstracts. *IBM Journal of research and development*, 2(2):159–165.
- A. Mandal, Paheli Bhattacharya, Sekhar Mandal, and Saptarshi Ghosh. 2021. Improving legal case summarization using document-specific catchphrases. In *International Conference on Legal Knowledge and Information Systems*.
- William C Mann and Sandra A Thompson. 1988. Rhetorical structure theory: Toward a functional theory of text organization. *Text-interdisciplinary Journal for the Study of Discourse*, 8(3):243–281.
- Kathleen McKeown and Dragomir R Radev. 1995. Generating summaries of multiple news articles. In *Proceedings of the 18th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 74–82.
- Kathleen R McKeown. 1985. Discourse strategies for generating natural-language text. *Artificial intelligence*, 27(1):1–41.

- Chris Mellish, Alistair Knott, Jon Oberlander, and Mick O'Donnell. 1998. Experiments using stochastic search for text planning. In *Proceedings of the 9th International General Workshop*, pages 98–107. ACL Anthology.
- Derek Miller. 2019. Leveraging bert for extractive text summarization on lectures. *arXiv preprint arXiv:1906.04165*.
- Gianluca Moro and Luca Ragazzi. 2022. Semantic self-segmentation for abstractive summarization of long documents in low-resource regimes. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pages 11085–11093.
- Amit Moryossef, Yoav Goldberg, and Ido Dagan. 2019. Step-by-step: Separating planning from realization in neural data-to-text generation. *arXiv preprint arXiv:1904.03396*.
- Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018. Ranking sentences for extractive summarization with reinforcement learning. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1747–1759.
- Shashi Narayan, Joshua Maynez, Jakub Adamek, Daniele Pighin, Blaz Bratanić, and Ryan McDonald. 2020. Stepwise extractive summarization and planning with structured transformers. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4143–4159.
- Shashi Narayan, Joshua Maynez, Reinald Kim Amplayo, Kuzman Ganchev, Annie Louis, Fantine Huot, Anders Sandholm, Dipanjan Das, and Mirella Lapata. 2023. Conditional generation with a question-answering blueprint. *Transactions of the Association for Computational Linguistics*, 11:974–996.
- Shashi Narayan, Yao Zhao, Joshua Maynez, Gonçalo Simões, Vitaly Nikolaev, and Ryan McDonald. 2021. Planning with learned entity prompts for abstractive summarization. *Transactions of the Association for Computational Linguistics*, 9:1475–1492.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pages 1532–1543.
- Laura Perez-Beltrachini and Mirella Lapata. 2018. Bootstrapping generators from noisy data. *arXiv preprint arXiv:1804.06385*.
- Seth Polsley, Pooja Jhunjhunwala, and Ruihong Huang. 2016. Casesummarizer: a system for automated summarization of legal texts. In *Proceedings of COLING 2016, the 26th international conference on Computational Linguistics: System Demonstrations*, pages 258–262.
- Ratish Puduppully, Li Dong, and Mirella Lapata. 2019. Data-to-text generation with entity modeling. *arXiv preprint arXiv:1906.03221*.
- Ratish Puduppully, Yao Fu, and Mirella Lapata. 2022. Data-to-text generation with variational sequential planning. *Transactions of the Association for Computational Linguistics*, 10:697–715.
- Ratish Puduppully and Mirella Lapata. 2021. Data-to-text generation with macro planning. *Transactions of the Association for Computational Linguistics*, 9:510–527.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 21(1):5485–5551.
- Hannah Rashkin, Maarten Sap, Emily Allaway, Noah A Smith, and Yejin Choi. 2018. Event2mind: Commonsense inference on events, intents, and reactions. *arXiv preprint arXiv:1805.06939*.
- Oliver Ray, Amy Conroy, and Rozano Imansyah. 2020. Summarisation with majority opinion. In *JURIX*, pages 247–250.
- Ehud Reiter and Robert Dale. 1997. Building applied natural language generation systems. *Natural Language Engineering*, 3(1):57–87.
- TYS Santosh, Vatsal Venkatkrishna, Saptarshi Ghosh, and Matthias Grabmair. 2024a. Beyond borders: Investigating cross-jurisdiction transfer in legal case summarization. *arXiv preprint arXiv:2403.19317*.
- TYSS Santosh, Cornelius Weiss, and Matthias Grabmair. 2024b. Lexsumm and lext5: Benchmarking and modeling legal summarization tasks in english. *arXiv preprint arXiv:2410.09527*.
- Murali Saravanan, Balaraman Ravindran, and Shivani Raman. 2006. Improving legal document summarization using graphical models. *Frontiers in Artificial Intelligence and Applications*, 152:51.
- Marijn Schraagen, Floris Bex, Nick Van De Luijngaarden, and Daniël Prijs. 2022. Abstractive summarization of dutch court verdicts using sequence-to-sequence models. In *Proceedings of the Natural Legal Language Processing Workshop 2022*, pages 76–87.
- Shirong Shen, Guilin Qi, Zhen Li, Sheng Bi, and Lusheng Wang. 2020. Hierarchical chinese legal event extraction via pedal attention mechanism. In *Proceedings of the 28th international conference on computational linguistics*, pages 100–113.
- Zejiang Shen, Kyle Lo, Lauren Yu, Nathan Dahlberg, Margo Schlanger, and Doug Downey. 2022. Multi-lexsum: Real-world summaries of civil rights lawsuits at multiple granularities. *Advances in Neural Information Processing Systems*, 35:13158–13173.

- Abhay Shukla, Paheli Bhattacharya, Soham Poddar, Rajdeep Mukherjee, Kripabandhu Ghosh, Pawan Goyal, and Saptarshi Ghosh. 2022. Legal case document summarization: Extractive and abstractive methods and their evaluation. In *Proceedings of the 2nd Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 12th International Joint Conference on Natural Language Processing*, pages 1048–1064.
- Santosh Tyss, Mahmoud Aly, and Matthias Grabmair. 2024. Lexabsumm: Aspect-based summarization of legal decisions. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 10422–10431.
- Ronald J Williams. 1992. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8:229–256.
- Wen Xiao, Iz Beltagy, Giuseppe Carenini, and Arman Cohan. 2022. Primera: Pyramid-based masked sentence pre-training for multi-document summarization. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5245–5263.
- Huihui Xu and Kevin Ashley. 2022. Multi-granularity argument mining in legal texts. *arXiv preprint arXiv:2210.09472*.
- Huihui Xu, Jaromir Savelka, and Kevin D Ashley. 2021a. Accounting for sentence position and legal domain sentence embedding in learning to classify case sentences. In *Legal Knowledge and Information Systems*, pages 33–42. IOS Press.
- Huihui Xu, Jaromir Savelka, and Kevin D Ashley. 2021b. Toward summarizing case decisions via extracting argument issues, reasons, and conclusions. In *Proceedings of the eighteenth international conference on artificial intelligence and law*, pages 250–254.
- Hui Yang, Tat-Seng Chua, Shuguang Wang, and Chun-Keat Koh. 2003. Structured use of external knowledge for event-based open domain question answering. In *Proceedings of the 26th annual international ACM SIGIR conference on Research and development in informaion retrieval*, pages 33–40.
- Feng Yao, Chaojun Xiao, Xiaozhi Wang, Zhiyuan Liu, Lei Hou, Cunchao Tu, Juanzi Li, Yun Liu, Weixing Shen, and Maosong Sun. 2022. Leven: A large-scale chinese legal event detection dataset. *arXiv preprint arXiv:2203.08556*.
- Yuheng Zha, Yichi Yang, Ruichen Li, and Zhiting Hu. 2023. Alignscore: Evaluating factual consistency with a unified alignment function. *arXiv preprint arXiv:2305.16739*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR.
- Shiyue Zhang, David Wan, and Mohit Bansal. 2022. Extractive is not faithful: An investigation of broad unfaithfulness problems in extractive summarization. *arXiv preprint arXiv:2209.03549*.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q Weinberger, and Yoav Artzi. 2019a. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*.
- Xingxing Zhang, Furu Wei, and Ming Zhou. 2019b. Hibert: Document level pre-training of hierarchical bidirectional transformers for document summarization. *arXiv preprint arXiv:1905.06566*.
- Linwu Zhong, Ziyi Zhong, Zinian Zhao, Siyuan Wang, Kevin D Ashley, and Matthias Grabmair. 2019. Automatic summarization of legal decisions using iterative masking of predictive sentences. In *Proceedings of the seventeenth international conference on artificial intelligence and law*, pages 163–172.
- Ming Zhong, Yang Liu, Da Yin, Yuning Mao, Yizhu Jiao, Pengfei Liu, Chenguang Zhu, Heng Ji, and Jiawei Han. 2022. Towards a unified multi-dimensional evaluator for text generation. *arXiv preprint arXiv:2210.07197*.

A Baselines

LED (Beltagy et al., 2020) is a longformer variant equipped with both encoder and decoder. In the encoder, it uses efficient local+global attention pattern instead of full self-attention, while the decoder utilizes full quadratic attention. It is initialized from the pre-trained BART model (Lewis et al., 2020). We use the LED-base version, which can handle input lengths of up to 16,384 tokens.

PRIMERA (Xiao et al., 2022) is initialized with LED model and pre-trained with a novel summarization-specific entity-based sentence masking objective, similar to Pegasus, but can be applied for longer inputs, motivated by Gap Sentence Generation objective of Pegasus (Zhang et al., 2020). It can handle upto 4096 tokens.

LongT5 (Guo et al., 2022) scales the input length by a new attention mechanism called as transient global attention which mimics local+global attention from ETC (Ainslie et al., 2020) and adopt summarization specific pre-training from PEGASUS into the T5 model (Raffel et al., 2020) for longer sequences. We use the LongT5-base version which can handle upto 16384 tokens.

SLED (Ivgi et al., 2023) processes long sequences by partitioning the input into overlapping chunks

and encode each chunk with a short-range pre-trained models encoder. It relies on decoder to fuse information across chunks attending to all input tokens, similar to fusion-in-decoder (?). SLED can be applied on top of any short-range models and thus we derive SLED-BART-base by applying on BART-base (Lewis et al., 2020).

Unlimiformer (Bertsch et al., 2023) adopts a strategy similar to SLED, but rather than attending to all input tokens, it focuses exclusively on the top-k tokens retrieved from a k-nearest-neighbor index constructed over the hidden states of all input tokens at every decoder layer. This helps to process unbounded length during testing, a key differentiator from SLED, which is constrained by memory when attending to all input tokens in the decoder. We derive Unlimiformer-BART-base model.

B Implementation Details

For the content selection selection, we created high-ROUGE episodes for the training set as described in original paper of (Gu et al., 2022) with a branching size of 2. We select the highest-scoring model based on the validation ROUGE-L scores. During inference, we set the stopping threshold to 0.6 and the maximum extracted sentences as 45 sentences. We train the model end-to-end for 20 epochs with a batch size of 1 using Adam Optimizer (Kingma and Ba, 2014) with β_{11} and β_{22} set to 0.9 and 0.999. We choose learning rate to be $1e-4$. During training, local sentence embeddings are kept fixed using pretrained Glove word embeddings of dimension $d = 200$ (Pennington et al., 2014). We set number of Bi-LSTM layers in local sentence and global context encoder to 2. For, Extraction History Encoder, we use 3 attention layers with 8 attention heads and a feed-forward hidden layer of a dimension of 1024 and dropout rate of 0.1. We run the Plumber framework on the summaries of all datasets to obtain the golden plans and prepare the plan generation datasets. We trained the plan generation unlimiformer models for 10 epochs on a batch size of 1. We set the maximum input length to 16384 tokens and the maximum generation length to 512 tokens. We set the chunk size and chunk overlap to 256 and 0.5 respectively. The model is trained end-to-end with learning rate of 0.0001 with Adam optimizer. Best model is chosen based on validation score of ROUGE-L between generated and golden plan. For summarization model, we trained LED for 15 epochs on a batch

size of 1 and with a learning rate of $3e-5$ using Adam Optimizer. The maximum input length is set to 16384 tokens and the maximum output length to 1024 tokens and the attention window size is set to 256 tokens. We chose the best model based on validation ROUGE-L score.

Plan A:

Two third-party nominees - filed - case | case - filed - in light of public health risks of COVID-19 | requiring signatures - violated - the First and Fourteenth Amendments | Court - ordered - defendants to reduce signature requirements | plaintiffs - sought -relief | defendants - reduce - signatures requirements by 30%

Summary A:

Two third-party nominees for the Georgia State House of Representatives filed this case in light of public health risks of COVID-19, alleging that the requirement to gather more than 20,000 signatures to qualify for the ballot violated their First and Fourteenth Amendment rights. The plaintiffs sought declaratory and injunctive relief to which the court ordered the defendants to reduce the required signatures by 30% ...

Plan B:

Two third-party nominees - filed - case | lawsuit - filed - in light of public health risks of COVID-19 | **signatures - required from- 5% of the registered voters** | requiring signatures - violated - the First and Fourteenth Amendments | Court - ordered - defendants to reduce signature requirements | **plaintiffs sought relief** | defendants - reduce - signatures requirements by 30%

Summary B:

Two third-party nominees for the Georgia State House of Representatives filed this case in light of public health risks of COVID-19, alleging that the requirement to **signatures from 5% of the registered voters in the district where they are running**, to qualify for the ballot violated their First and Fourteenth Amendment rights. The court ordered the defendants to reduce the required signatures by 30% ...

Table 7: Case Study illustrating how plans offer controllability to generate customised summaries.

C Case Study

Table 8 illustrates a summary generated using a baseline model without planning, alongside models that incorporate entity and event-centric planning. This case study demonstrates how event-centric planning can enhance both the structural coherence and the faithfulness of a summary to the source document.

Table 7 showcases the customization process of modifying plans, which can influence the final summary. This case study illustrates how event-centric planning offers users the flexibility to create tailored summaries.

Reference Summary

... The Defendants were Furniture Distributors, Inc., Kimbrell's of Statesville, NC, and Kimbrell's, Inc., all associated as a furniture store. The controversy involved the alleged sexual harassment and constructive discharge of an employee in violation of Title VII of the Civil Rights Act of 1964. The lawsuit was filed in August of 2002 and was eventually resolved by consent decree in October of 2003. The complainant, who intervened in the case, was awarded \$25,000 and a guarantee of neutral references. ...

w/o Planning

... against Furniture Distributors, Inc., Kimbrell's Furniture, Inc., and Kimbrell's, Inc. in the U.S. District Court for the Western District of North Carolina alleging discrimination on the basis of sex in violation of Title VII of the Civil Rights Act of 1964. Specifically, the complaint alleged that the defendant subjected the charging party, a female employee, to a sexually hostile work environment and constructively discharged her. Following some discovery, **the parties settled the lawsuit in August 2002 through a consent decree**. The three-year decree, containing non-discrimination and non-retaliation clauses, required the defendant to: report on complaints and make other compliance reports, allow the EEOC access for monitoring ...

Planning with Events

Plan:

... Furniture Distributors, Inc. | Kimbrell's Furniture | Statesville | NC | Kimbrell's ||| August of 2002 | October of 2003 ||| 25,000 ||| ...

Summary:

... The defendant was Furniture Distributors, Inc., a furniture manufacturer, doing business as Kimbrell's Furniture, a general store in Statesville, NC, and Kimbrell's, Incorporated collectively doing business as Kimbrell's. The complaint alleged that the defendant subjected the charging party to a sexually hostile work environment and constructively discharged her. **The charging party intervened in the lawsuit in August of 2002** and in October of 2003 the parties settled. The consent decree required the defendant to pay the charging party \$25,000. ...

CoPERLex: Planning with Events

Plan:

... Defendants - were - Furniture Distributors | controversy - involved - sexual harassment of employee | sexual harassment - violation - Title VII of the Civil Rights Act of 1964 | Lawsuit - filed - August 2002 | Lawsuit - resolved - consent decree | Resolution - occurred - October 2003 | Complainant- awarded - \$25,000 ...

Summary:

... The Defendants were Furniture Distributors, Inc., Kimbrell's Furniture, Inc., and Kimbrell's, Inc., all associated as a furniture store. The controversy involved sexual harassment of an employee in violation of Title VII of the Civil Rights Act of 1964. The lawsuit was filed in August of 2002 and was eventually resolved by consent decree in October of 2003. The complainant was awarded \$25,000.00. ...

Table 8: Case study demonstrating the enhancement of structural coherence and faithfulness to source documents through event-centric planning.