

Can LLMs Time Travel? Enhancing Temporal Consistency in Legal Agentic Search through Reinforcement Learning

Wei Fan¹, Yining Zhou¹, Mufan Zhang¹, Yanbing Weng¹, Yiran Hu³, Tianshi Zheng¹,
Baixuan Xu¹, Chunyang Li¹, Jianhui Yang², Haoran Li^{1*}, Yangqiu Song¹

¹Department of Computer Science and Engineering, HKUST, Hong Kong SAR, China

²School of Law, Tsinghua University, Beijing, China

³Cheriton School of Computer Science, University of Waterloo, Waterloo, Canada

wfanag@connect.ust.hk, hlibt@connect.ust.hk, yqsong@cse.ust.hk

Abstract

While large language models (LLMs) augmented with agentic search capabilities show promise for legal reasoning, they overlook a fundamental constraint that applicable law must match the temporal context of each case, as retroactive application of statutes violates core legal principles and leads to erroneous conclusions. Our observations reveal that current legal LLMs suffer from temporal bias anchored to their training cutoff, while search agents rarely incorporate temporal constraints into queries, and that web search alone cannot provide the precise statute and precedent citations that legal reasoning demands. To address these challenges, we propose **LegalSearch-R1**, an end-to-end reinforcement learning framework that pairs local statute RAG for precise article matching with online web search for broader legal knowledge, trained on temporally-indexed data spanning multiple amendment periods to enforce temporal consistency. Extensive experiments on our benchmark covering 13 legal tasks demonstrate that our 7B-parameter agent outperforms state-of-the-art deep research frameworks and specialized legal LLMs by 12.9% to 29.8%, surpasses baselines by 57.7% to 80.3% on temporal consistency, and exhibits robust out-of-domain generalization. The code and data are available at <https://github.com/HKUST-KnowComp/LegalSearch-R1>.

*Corresponding Authors