

GRACE: Grounded Adversarial Reasoning over Canadian Law

Jiakang Xu^{1*}, Wantong Huo^{1*}, Udom Silparcha², Jonathan H. Chan³

¹University of Toronto, ²UTOA Computing Analytics Co., Ltd.,

³King Mongkut’s University of Technology Thonburi

{amyj.xu, vanessa.huo}@mail.utoronto.ca, udom@utoa-thai.com, jonathan@sit.kmutt.ac.th

Abstract

Large language models have shown strong performance across a range of legal tasks, but existing benchmarks rarely evaluate the ability to take and defend a legal position, reason under incomplete information, or synthesize multiple statutory provisions. This gap is particularly pronounced for Canadian law, which remains underrepresented in legal NLP. We introduce **GRACE** (Grounded Reasoning Adversarial Canadian lEgal examples), a dataset of 1,915 question–reasoning–answer instances grounded in Canadian federal legislation. GRACE covers three reasoning modes: adversarial advocacy, uncertainty, and applied reasoning. We develop a pipeline that partitions raw statutory text, generates scenario-based questions and reasoning, and filters examples through model-free citation verification and LLM-based quality auditing. As a proof of concept, we fine-tune **CLeAR-4B** (Canadian Legal Adversarial Reasoning), a lightweight model for grounded legal reasoning, and evaluate it against the unmodified Qwen3-4B base model in open- and closed-book settings. CLeAR-4B substantially improves agreement with teacher outputs and statutory citation behavior when the relevant act text is provided, while its grounding degrades sharply when the statute is withheld. These results suggest that GRACE can support the development of lightweight legal models that reason more effectively from supplied statutory text.¹

*Equal contribution.

¹Data and code:

<https://github.com/amyjiakangxu-code/GRACE-CLeAR>