

Grant Writing Sample

SR&ED Technical Narrative: Adaptive Compliance Monitoring Tool

Portfolio Piece | Nabila Husseni | nabilahusseni.com

About This Sample

This is a writing sample created for portfolio purposes. It shows how I write SR&ED (Scientific Research and Experimental Development) grant applications for Canadian technology companies. All company names and project details are fictional. The structure, language, and format follow real SR&ED submission standards as set by the Canada Revenue Agency.

Project Overview

Project Title:

Building a Smarter Compliance Monitoring Tool Using Machine Learning

Company: TechNova Solutions Inc. (fictional) | Fiscal Year: January to December 2024

TechNova Solutions builds compliance software for mid-sized Canadian financial institutions. When regulators update their rules, TechNova's clients need their systems updated quickly. The old way of doing this involved manual rule changes, specialist support, and weeks of testing. This project explored whether a machine learning model could read new regulatory documents, understand what changed, and suggest system updates automatically, so clients could respond to regulatory changes in days instead of weeks.

The Problem We Were Trying to Solve

Canadian financial regulators (OSFI, FINTRAC, CDIC) publish updates regularly. Each update uses highly technical legal language that is full of cross-references and conditional rules. Standard language processing tools, trained on everyday text, could not read this kind of writing accurately. They confused guidance documents with binding rules, and missed the difference between a rule that always applies and one that only applies in certain situations.

TechNova's compliance team was spending significant time manually interpreting new rules and translating them into system logic. The company needed to know: could a machine learning model be trained specifically on Canadian regulatory documents to do this work faster and more reliably?

What Was Uncertain at the Start

Three things were genuinely unknown when the project began:

First, no one knew whether there was enough labeled Canadian regulatory data available to properly train a specialized model. General-purpose language models are trained on millions of documents. Regulatory documents from Canadian financial agencies are a small, specialized category with very little pre-existing labeled training data.

Second, even if the model could classify regulatory requirements accurately, it was unclear whether it could assign confidence scores that would be useful to human reviewers. A tool that flags everything as high-risk shifts the burden without solving the problem.

Third, the team did not know whether the model could process new regulatory documents fast enough to meet the client requirement of generating a proposed system update within 48 hours of a regulatory publication.

What the Project Achieved

The team built a training dataset by collecting 14 years of Canadian financial regulatory documents and having compliance specialists annotate them. This took longer than expected and required several rounds of review to reach acceptable quality. The result was a training corpus that did not previously exist in this form.

A fine-tuned language model trained on this dataset improved classification accuracy on conditional regulatory rules by 34 percent compared to a general-purpose baseline model. This was a meaningful improvement in a domain where misclassification has compliance and legal consequences.

In pilot testing with two enterprise clients, the average time from regulatory publication to a validated system update dropped from 22 days to 4 days. This exceeded the 48-hour target for initial proposal generation while meeting it for analyst-reviewed implementation. Both pilot clients confirmed the tool was usable and accurate enough to reduce their manual review workload significantly.

How the Work Was Done

The project ran for 12 months and followed a structured experimental process. The team began by reviewing available language model architectures and documenting the reasons for selecting the approach they chose. They then developed annotation guidelines, tested them across multiple reviewers for consistency, and revised them before applying them to the full training corpus.

Model training ran across six experimental cycles. Each cycle started with a hypothesis, produced a trained model, and ended with a documented analysis of where the model succeeded and where it failed. Three of the six cycles required the team to revise their assumptions and start the next cycle differently. This iterative process is a core feature of SR&ED-eligible work.

Pilot deployment with two enterprise clients provided real-world testing data. The team identified two edge cases in the confidence scoring logic that required additional experimental work before the model was ready for broader use.

Why This Qualifies as SR&ED

This project meets the three CRA criteria for SR&ED eligibility. There was genuine scientific uncertainty: the team did not know in advance whether a domain-specific model trained on Canadian regulatory text would perform well enough to be useful. There was scientific advancement: the annotated training corpus and fine-tuning methodology represent new knowledge that did not previously exist in this domain. And the work was conducted through systematic investigation: each experimental cycle was hypothesis-driven, documented, and informed the next cycle.

This sample was created for portfolio purposes. All company names, client references, and project details are fictional. Structure and language reflect real SR&ED submission requirements. | nabilahusseni.com