

INTRODUCTION

This project focuses on classifying SNOMED CT concept codes¹ within Singapore's **National Electronic Health Record (NEHR)** as **Sensitive Health Information (SHI)** – conditions carrying heightened risk of stigma or discrimination, which are placed under stricter access safeguards. These codes are tagged across ten SHI categories²:

1. Sexually Transmitted Infections (STIs),
2. Human Immunodeficiency Virus (HIV) Infection / Acquired Immunodeficiency Syndrome (AIDS),
3. Schizophrenia & Delusional Disorder,
4. Substance Abuse & Addictions,
5. Biological Parenthood,
6. Voluntary Sterilisation,
7. Termination of Pregnancy,
8. Organ Transplant,
9. Suicide or Attempted Suicide, and
10. Domestic, Child or Sexual Abuse.

Classification is currently **manual, blind to semantic variation** beyond curated keywords, and **slow to track amendments** to the Misuse of Drugs Act (MDA).

OBJECTIVES

1. **Streamline the workflow** – surface and classify concepts end-to-end.
2. **Enhance semantic awareness & regulatory adaptability** – replace curated criteria with learned embeddings; track MDA updates live.
3. **Optimise manpower** – auto-accept confident cases, route only ambiguous ones to experts.

DATASET

1. Classifier Dataset → Component 1

Trains and evaluates the semantic classifier across categories. *Schizophrenia & Delusional Disorder* excluded (handled separately).

- **Split:** 80/10/10 (Train/Val/Test) at **concept level** (prevents data leakage)
- **Training set:** Downsampled to **50×** through experimentation as Non-sensitive >> sensitive.
- **Val/Test set:** Original distribution retained

2. Substance Extractor Dataset → Component 2

Evaluates the generative substance extractor.

Total 1,500 concepts, covering the following substance space:

- **Non-abusive substances**
- **Controlled substances**
- **Broad SNOMED CT concepts** – mostly substance-free

3. Agentic LLM Evaluation Dataset → Component 3

Evaluates agentic LLM end-to-end on unseen concepts.


- Derived from the **classifier test set**
- Non-sensitive downsampled due to resource constraints

REFERENCES

1. SNOMED International [Accessed: 2026-07-24] <https://browser.ihtsdo.org/>
2. Ministry of Health. (2025). *HealthTech Instruction Manual (HiIM)*. <https://hpp.moh.gov.sg/healthtech-instruction-manual>

PIPELINE ARCHITECTURE

Pipeline consists of two upstream signal generators feeding an agentic decision-maker:



RESULTS & ANALYSIS

Component 1: Semantic Classifier Module

1. Base Model Selection


BERT, PubMedBERT and SapBERT compared on final-layer embeddings with a logistic regression head. **PubMedBERT** selected (macro F1 0.74).

2. Pretrained Embeddings → Fine-Tuning

Frozen pretrained embeddings tried across logistic regression (threshold-tuned), Random Forest, and XGBoost. However, no set up met the recall ≥ 0.90 bar. **Fine-tuning the encoder lifted recall across every classifier.**

Final model: logistic regression at threshold 0.8 on fine-tuned layers 10 & 11 chosen as the simplest configuration clearing the bar.


Component 2: Substance Extraction Module



Final chosen setup: Gemma 3 with a separate critic verifying and correcting over a feedback loop (up to 3 retries).

Component 3: Agentic LLM with Dynamic Retrieval

a. Runtime




Active time reduced from 199 → 8 minutes (96%), primarily by cutting the volume requiring human classification efforts.

b. Classification Performance

		Confusion matrix									
		Auto-decided subset (blue, n = 746)					Flagged subset (yellow, n = 42)				
ACTUAL	Non-sensitive (n = 229)	228	0	0	0	0	0	0	0	0	1
	Substance Abuse / Addiction (n = 147)	0	134	0	0	0	0	0	0	0	13
	STI (n = 83)	0	0	80	0	0	0	0	0	0	3
	Organ Transplant (n = 85)	1	0	0	78	0	0	0	0	0	6
	Termination of Pregnancy (n = 72)	0	0	0	0	68	0	0	0	0	4
	Suicide / Attempted Suicide (n = 63)	0	0	0	0	0	58	0	0	0	5
	HIV / AIDS (n = 58)	0	0	0	0	0	0	55	0	0	3
	Domestic / Child / Sexual Abuse (n = 28)	0	0	0	0	0	0	0	26	0	2
	Voluntary Sterilisation (n = 20)	0	0	0	0	0	0	0	0	16	4
	Biological Parenthood (n = 3)	0	0	0	0	0	0	0	0	0	1
	Non-sensitive										
	Substance Abuse / Addiction										
	STI										
	Organ Transplant										
	Termination of Pregnancy										
	Suicide / Attempted Suicide										
	HIV / AIDS										
	Domestic / Child / Sexual Abuse										
	Voluntary Sterilisation										
	Biological Parenthood										
	Flagged for human review										
		PREDICTED									

~5.1% of cases flagged for human review, with reasoning trace provided to support decision-making process by human; **0.999 recall on the remaining 95% of cases.**

c. Explanation Quality



Every trace is at minimum defensible – **73–95%** fully grounded per category, and **83%** among flagged concepts,

DISCUSSION

The pipeline delivers on all three objectives:

1. **Streamlined Workflow:** active time falls **96.0%**; criteria upkeep and statute-checking are eliminated, and classification is confined to the flagged subset.
2. **Semantic Awareness & Regulatory Adaptability:** Live MDA First Schedule retrieval keeps classification aligned with the latest statute. Embedding-based scoring gives every concept a scored, auditable prediction.
3. **Optimised Manpower:** Experts review **5.3%** rather than the whole set, and validate against a reasoning trace rather than classifying from scratch.