أكاديمية كاوست
KAUST ACADEMY

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

UNIVERSITY OF TORONTO

Analyzing Social Determinants of Health in Ontario Hospital Feedback: A Heart Failure Patient Case Study

Name	Contribution
Rama Alyoubi	Identifying SDoH in patient comments
Minh Le	Applying heart failure use case

Supervisor
Prof. Zahra Shakeri

Data Science Institute, University of Toronto

Aug 27, 2025

أكاديمية كاوست
KAUST ACADEMY

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

Abstract

Ontario hospitals collect thousands of post-discharge surveys, yet until a joint investigation by the *Toronto Star* and the Investigative Journalism Bureau compelled their release, patient feedback remained inaccessible. Our analysis of **120,695 anonymized comments from 47 hospitals (2015–2022)** revealed a strained system: nearly 80% of facilities received more negative than positive feedback, and patients frequently described how poor care experiences and systemic inequities eroded their health, dignity, and trust.

Social Determinants of Health (SDoH) including income, housing, and social support—are increasingly recognized as critical drivers of outcomes. However, these factors have rarely been extracted from large-scale free-text data due to the difficulty of labeling and aggregating across institutions. In this work, we successfully developed a **scalable pipeline** to detect and analyze SDoH signals at scale. Our case study focused on **cardiovascular disease, Ontario's leading cause of death**, whose outcomes are tightly linked to SDoH. We fine-tuned large language models (LLMs) to identify heart-failure patient comments and combined this with SDoH detection to quantify equity barriers specific to this vulnerable group. Partnering with **Transform HF**, we refined our analysis through patient co-design and integrated the results into an **interactive dashboard**.

From the dataset, we detected **16,132 comments (13.5%)** containing SDoH, with **Living Condition (21%)** and **Financial Constraint (19%)** emerging as the most frequent determinants. Within the **5,133 heart-patient comments**, **924 SDoH instances (18%)** were found, again dominated by Living Condition and Financial Constraint, underscoring their consistent influence across both the general dataset and the heart-failure subgroup. Model agreement analysis showed that Gemma and LLaMA aligned most closely, while Phi-3 diverged more significantly. To address this variability, we applied majority voting across all four models, which reduced single-model biases while preserving sensitivity to less frequent categories.

Our deliverables include a **fully labeled dataset**, **fine-tuned checkpoints** for SDoH tagging and heart-patient identification, and a **working dashboard** that enables province-wide and subgroup-specific insights. By fusing natural language processing, domain adaptation, and participatory design, this project transforms hidden narratives into actionable knowledge and contributes new insights toward closing Ontario's health equity gap.

Table of Contents

Project Objectives and Goals....4
Current Progress and Milestones Achieved....4
Methodology and Approach....4
Preliminary Results or Findings....6
Challenges and Obstacles Encountered....7
Next Steps and Future Work....7
Resources and Tools Used....7

أكاديمية كاوست
KAUST ACADEMY

جامعة الملك عبدالله للعلوم والتقنية
King Abdullah University of Science and Technology

Project Objectives and Goals

To build an AI-powered tool that highlights how social factors affect patient care, starting with heart failure patients.

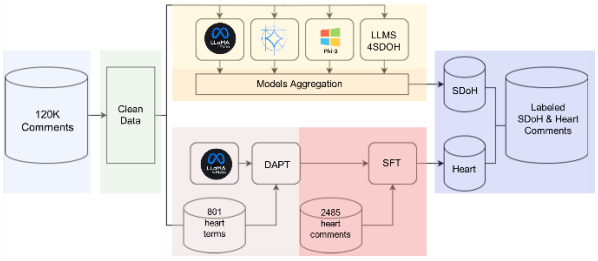
Project Objectives:

- Detect Social Determinants of Health (SDoH) from Ontario hospital patient feedback using large language models (LLMs).
- Apply a **heart failure case study** to examine how SDoH specifically impact cardiovascular outcomes.
- Build an **interactive dashboard** to visualize findings and engage stakeholders (e.g., Transform HF) in co-design.

Current Progress and Milestones Achieved

- ✓ Cleaned a dataset of **120,695 patient comments** from 47 Ontario hospitals (2015–2022).
- ✓ Removed **6,325 duplicates and empty comments** during preprocessing.
- ✓ Applied LLM-based labeling across **11 SDoH categories**.
- ✓ Fine-tuned models to identify **2,485 heart-related comments**, yielding **5,133 heart patient records** with both SDoH and heart labels.
- ✓ Built an interactive dashboard for analysis.

Methodology and Approach



Our methodology is structured as a multi-stage pipeline that processes patient feedback, detects Social Determinants of Health (SDoH), and identifies heart-patient subgroups for deeper analysis.

1. Data Preparation

a. Collected 120,695 patient comments from 47 Ontario hospitals (2015–2022).

b. Cleaned the dataset by removing 6,325 duplicate and empty entries and applying tokenization and normalization to prepare the text for modeling.

2. SDoH Detection via Model Aggregation

Model	Developer	Parameters	Context
Gemma-2B-it	DeepMind	2 billion	4K tokens
Phi-3-it	Microsoft	3.8 billion	128K tokens
Llama-3-8B-it	Meta	8 billion	8K tokens
LLMs4SDoH	Yale	7 billion	4K tokens

a. Applied four large language models (LLMs): LLaMA-3, Gemma-2B-it, Phi-3-it, and LLMs4SDoH.

b. Each model classified patient comments into 11 SDoH categories (adapted from Yu et al., 2024).

c. Used majority voting to aggregate outputs, ensuring robust labeling.

d. Evaluated agreement between models using Hamming distance and Cohen's Kappa, which highlighted stronger alignment between Gemma and LLaMA compared to Phi-3.

3. Heart Patient Identification

- Developed a two-stage pipeline to isolate comments from heart-failure patients:
    - Domain Adaptive Pre-Training (DAPT): Adapted the LLaMA model using 801 heart-related terms combined with the full dataset to enhance domain knowledge.
    - Supervised Fine-Tuning (SFT): Applied to 2,485 heart-related comments, producing a model optimized for identifying heart-patient feedback.
- Training was conducted using Low-Rank Adaptation (LoRA) on a single L40 GPU.

4. Integration of SDoH and Heart Labels

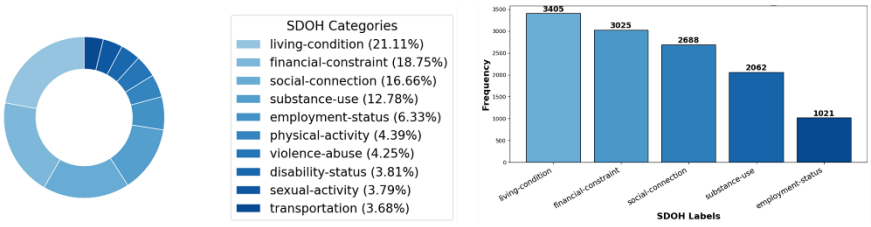
a. Combined both outputs (SDoH labels from model aggregation and heart labels from fine-tuning) into a unified dataset.

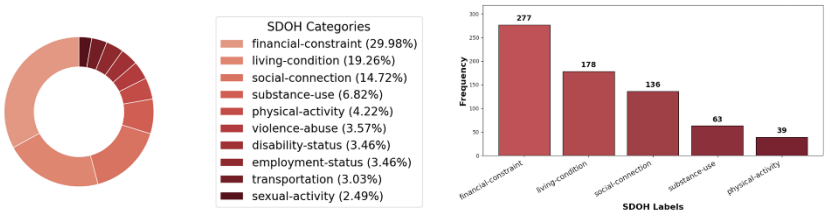
b. Produced 16,132 SDoH-labeled comments and 5,133 heart-patient comments, including 924 heart-patient comments with SDoH detected.

c. Final dataset supports visualization through the interactive dashboard, enabling co-design with Transform HF and stakeholders for actionable insights.

Preliminary Results or Findings

- From **120,695 patient comments**, we detected **16,132 comments (13.5%)** containing Social Determinants of Health (SDoH).
- The most common determinants were **Living Condition (21%)** and **Financial Constraint (19%)**, reflecting housing conditions, food access, and affordability as major themes.
- Among **heart-patient comments (5,133 identified)**, **924 SDoH instances (18%)** were found. Living Condition and Financial Constraint again emerged as the most frequent categories, showing consistency across both the general dataset and the heart-failure subgroup.





- **Model agreement analysis** showed that **Gemma and LLaMA aligned most closely**, while pairs involving Phi-3 had larger differences. To address this variability, we applied **majority voting** across all four models, which reduced single-model biases while maintaining sensitivity to less frequent SDoH categories.

Model Pair	Hamming Distance	Cohen's Kappa
Gemma ↔ LLaMA	0.1366	0.2431
Gemma ↔ Phi3	0.3849	0.0358
LLaMA3 ↔ Phi3	0.3139	0.0631
Overall Average	0.2784	0.1140

Challenges and Obstacles Encountered

- Variability in model performance (Gemma and LLaMA aligned closely, while Phi-3 diverged significantly).
- Limited labeled data for certain SDoH categories (e.g., race, gender) led to exclusion or merging of categories.
- Large dataset size required significant compute resources and optimization for efficient training.

Next Steps and Future Work

- Refine dashboard visualizations and integrate co-design feedback from heart-failure patient partners.
- Extending analysis to **other chronic conditions** beyond cardiovascular disease.
- Improve multi-label classification robustness to capture overlapping SDoH categories.
- Document and finalize source code for reproducibility and public release.

Resources and Tools Used

- **Dataset:** 123,849 anonymized comments from 47 Ontario hospitals (2014–2022)
- **Compute Canada & KAUST Ibex clusters** for GPU-based training.
- **LLMs**: Gemma-2B-it, Phi-3-it, LLaMA-3-8B-it, and LLMs4SDoH.
- **Techniques**: Domain Adaptive Pre-Training (DAPT), Supervised Fine-Tuning (SFT), Low-Rank Adaptation (LoRA).