Entropy-Based Modeling For Detecting Behavioral Anomalies in Users of a Diabetes Lifestyle Management Support System

Identifying non-adherence indicators in a chatbot-based diabetes support system

Sorin-Andrei Ciuntu - s.a.ciuntu@student.tudelft.nl

Supervisors: C. M. Jonker, J. D. Top

1. Background

People living with diabetes face more rigorous demands compared to the general population when it comes to managing their health.

Regular glucose self-monitoring, careful diet control, and consistent physical activity are some of the most important lifestyle factors in managing this condition.

Tools such as CHIP, an AI-based chatbot prototype, exist to support lifestyle changes, yet many patients struggle to adhere to their treatment plans.

This study focuses on detecting behavioral shifts in an AI chatbot-based diabetes lifestyle management support system. Behavioral shifts have been shown to correlate with non-adherence inside medical contexts.

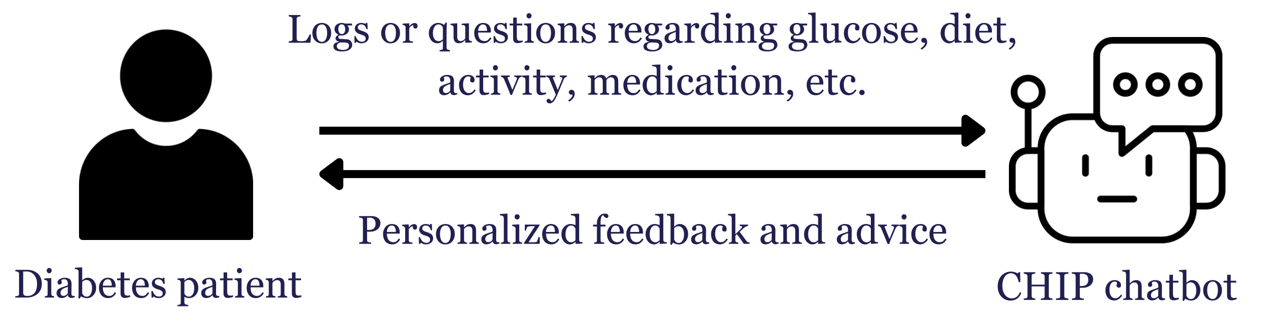

2. Research Question

How can entropy-based behavioral modeling help detect anomalies indicative of non-adherence in a chatbot-based diabetes lifestyle management support system?

3. Entropy

To quantify behavioral variability, we used Shannon entropy (Equation 1). This measures the randomness of a random variable. In our case, the empirical random variable is the user's observed behavior.

For a distribution **P** with outcomes $\mathbf{x_1}, \mathbf{x_2}, \ldots, \mathbf{x_n}$, the entropy is:

$$H(X) = -\sum_{i} P(x_i) \log P(x_i) \qquad (1)$$

To illustrate an example, a user with consistent logging at the same time each day has highly predictable behavior and therefore should have low baseline entropy. If the logging of this patient suddenly became erratic, the entropy would increase greatly above baseline.

3. System design

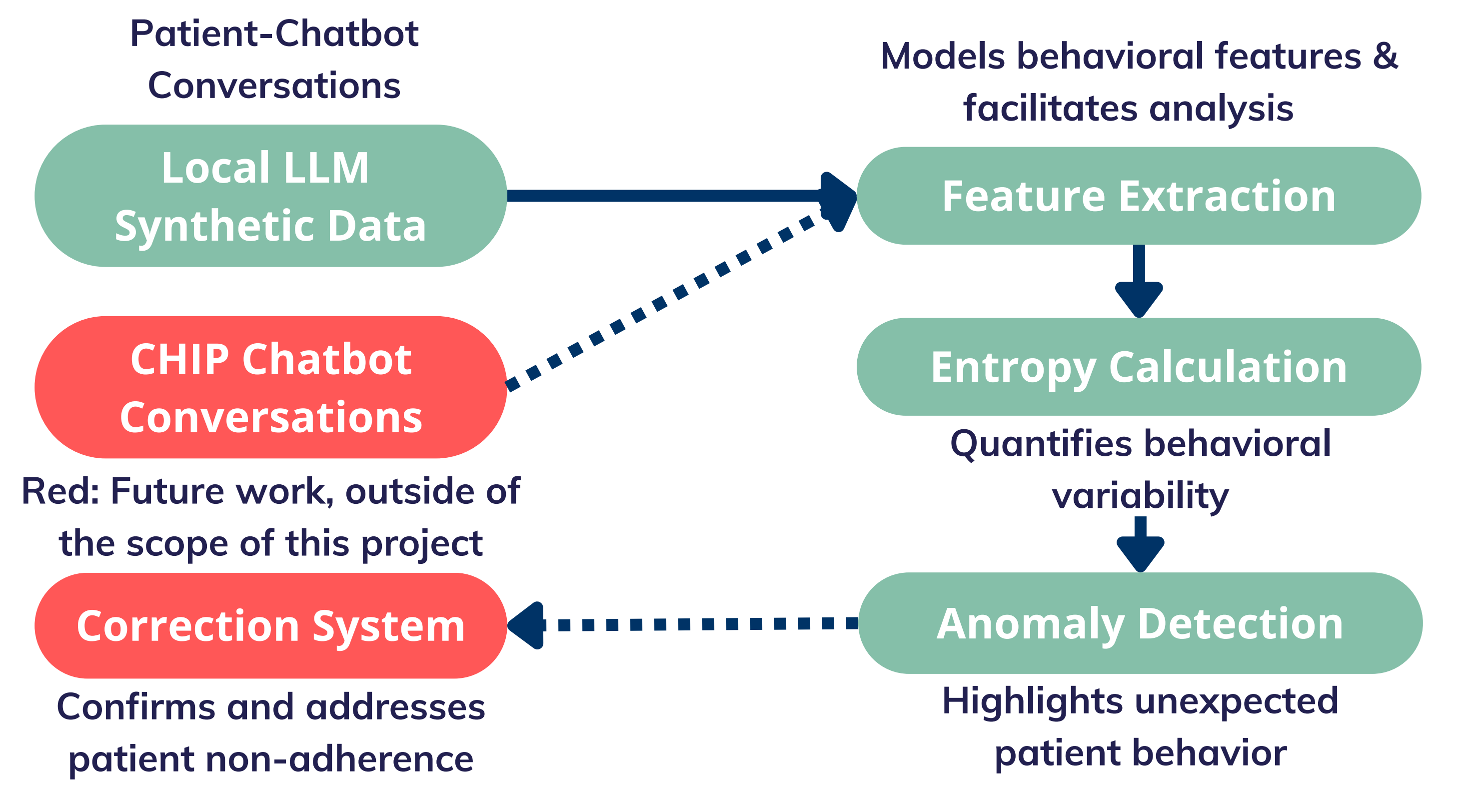

Synthetic patient-chatbot conversations were generated using a local language model, with some user profiles designed to change behavior midway. Entropy was computed over four features, and significant deviations flagged potential engagement or adherence shifts.

5. Results

As an example of the results obtained, the graph below shows the entropy evolution for the 'time of day' feature of the user persona who changes behavior from consistent, frequent, and varied content to erratic, infrequent, and similar content. Red dots indicate anomalies flagged by the detection system. The windows capturing the behavioral shift are highlighted in orange.

In this image, three marked anomalies fall within the highlighted window and are therefore true positives. The other two red points near the end represent false positives.

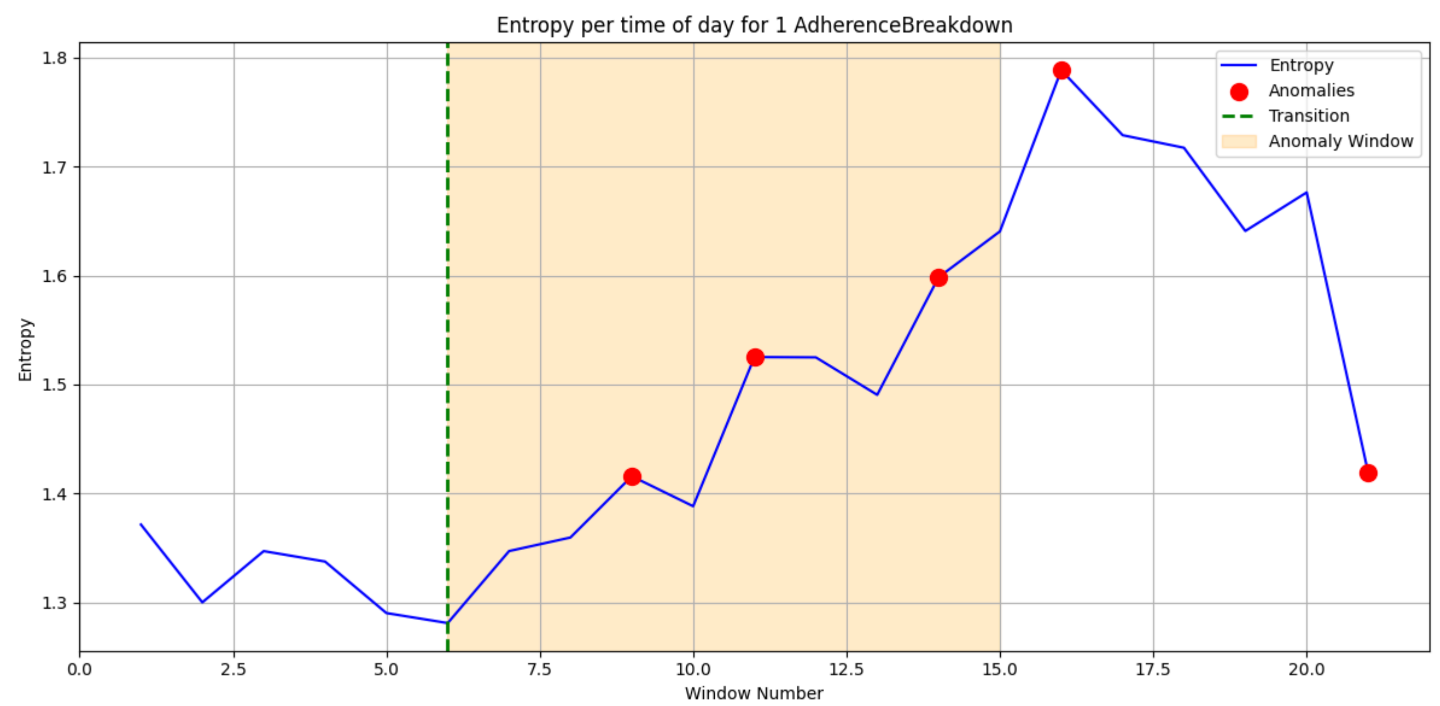

5. Results (cont.)

Running the system on the synthetic dataset yielded an overall accuracy of approximately 76%.

However, the false positive rate was relatively high at around 15%.

Also, the system correctly identified only around 35% of the actual anomalies.

The confusion matrix on the right summarizes the system's performance across all simulated user conversations.

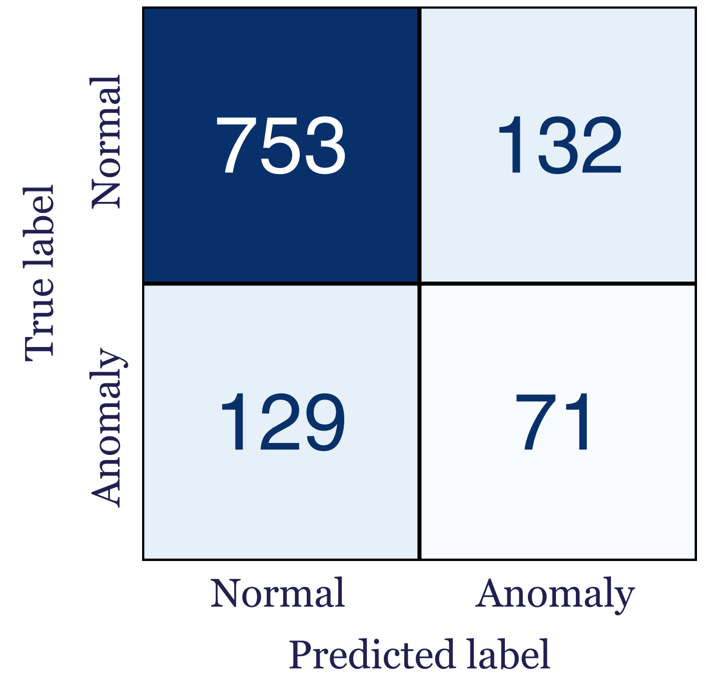

6. Conclusions

The results show that the system detects behavioral shifts in the experimental dataset with reasonable efficiency, but with much room to improve.

Notably, the system tends to over-flag normal behavior, especially in users with naturally high variability, as indicated by the high false positive rate of 15%. This rate should be reduced before the system can find practical use.

7. Future Work

To evaluate performance more realistically, future studies should use longitudinal data from real patients. This should include examining correlations between detected anomalies and clinical outcomes.

If validated, future research should focus on integrating this system into diabetes lifestyle management support systems like CHIP to enable early detection of patient disengagement or non-adherence and provide timely interventions for patients at risk.

Additionally, the current anomaly detection approach is simple and may need to be improved to increase accuracy for practical applications. We suggest exploring the following directions:

- **Methods for dynamic binning** such as equal-frequency binning or patient persona clustering
- **Personalized thresholds** derived from historical user data
- **Machine learning-based anomaly detection algorithms,** such as Isolation Forest or autoencoders