From Data to Daily Decisions: A Safety-Constrained, AI-Assisted N-of-1 Framework for Back Rehabilitation

Author: Hossein Mokhtarzadeh
Email: mokhtarzadeh.hossein@gmail.com
Website: poseiq.com

Abstract

This paper presents a lightweight, safety-constrained, semi-automated agentic decision-support framework for daily rehabilitation. The system combines structured user inputs, rule-based safety triage, and conservative plan generation to support consistent day-to-day decisions. In a single-case (N-of-1) implementation based on self-reported daily logs, the framework was associated with reduced symptom volatility, clearer identification of triggers, and improved standing and walking tolerance over the observation period. Here, “agentic” refers to a human-in-the-loop system that translates inputs into constrained, rule-based actions. The framework is non-diagnostic and is intended to support structured, safe decision-making under uncertainty.

1. Introduction

Rehabilitation outcomes are often limited not by lack of clinical knowledge, but by inconsistent day-to-day decisions, over-progression, and variable responses to symptom fluctuations. While biomechanical tools and clinical expertise provide strong guidance, the period between sessions is often inconsistently managed.

This work explores whether a lightweight, AI-assisted system can support these daily decisions through a structured, safety-constrained approach. The goal is to reduce variability, improve consistency, and complement existing clinical and biomechanical workflows.

This study aimed to evaluate the feasibility and practical utility of such a system in a single-case (N-of-1) implementation.

2. Framework Overview

The proposed system operates as a safety-constrained, semi-automated decision-support assistant. In this context, “agentic” refers to a human-in-the-loop system that translates structured user inputs into adaptive, rule-based actions (daily plans) under predefined safety and progression constraints.

The framework consists of three core components:

1. Structured daily check-ins (pain, symptoms, and functional tolerance)
2. Rule-based safety triage (RED / AMBER / GREEN)
3. Minimal, conservative daily plan generation

Operationally, inputs include pain_now, pain_worst, symptom location and spread, leg symptoms, sitting/standing/walking tolerance, triggers, relief factors, and red-flag screening. Triage is determined using predefined rules: RED indicates potential risk requiring escalation, AMBER indicates worsening or unstable symptoms requiring load reduction or holding, and GREEN indicates stable or improving status allowing cautious progression. Progression is constrained, with no advancement following RED or AMBER states and only small increases permitted under stable conditions.

The system is non-diagnostic and does not perform prediction, but focuses on structured, consistent, and safe decision-making.

3. Methods

This study used a single-case (N-of-1), longitudinal, practice-based design with daily logs collected from 2026-03-04 to 2026-03-20 (non-consecutive missing days: 2026-03-05, 2026-03-09, 2026-03-15).

A semi-automated, rule-based decision-support pipeline was applied daily:

Check-in → Safety triage (RED/AMBER/GREEN) → Short trend review → Conservative plan generation

Daily check-in fields included pain_now, pain_worst, symptom location and spread, leg symptom status, sitting/standing/walking tolerance, main trigger, main relief factor, and red-flag screening (Fig. 1).

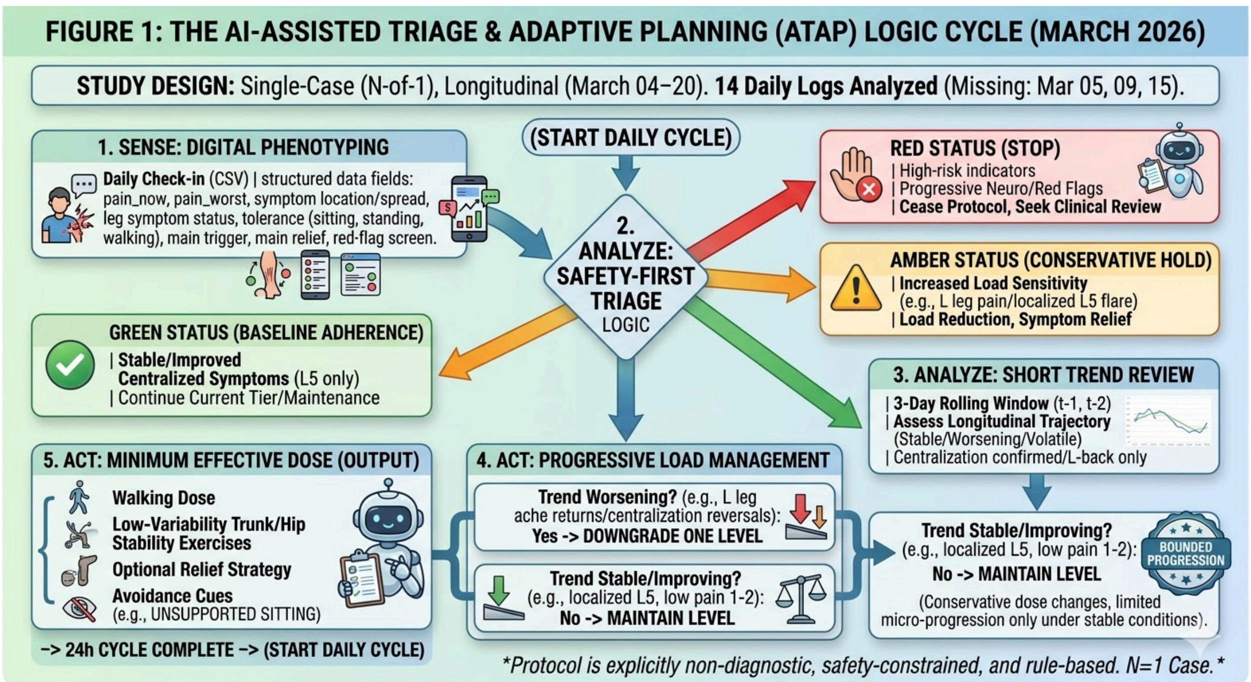

Figure 1: The AI-Assisted Triage and Adaptive Planning (ATAP) logic cycle. This schematic outlines the daily workflow used in the N-of-1 study, integrating structured daily self-reported inputs, rule-based safety triage (RED/AMBER/GREEN), and short-term trend review to generate a conservative, minimum-effective daily plan.

Trend review was based on recent logs (approximately a 72-hour window) to identify worsening, stable, or improving patterns. Symptom volatility was assessed descriptively as variation in reported pain levels across logs.

The system was explicitly safety-constrained and non-diagnostic. Progression was conservative and bounded, with small dose changes permitted only under stable conditions. No progression was allowed following RED or AMBER states.

Plan outputs were limited to:

- walking dose
- low-variability trunk/hip stability exercises
- optional relief strategy

- avoidance cues

Decision rules were predefined and applied consistently throughout the observation period.

Analysis was descriptive and based on available logs and generated plans; no inferential statistics were applied.

As a self-experiment (N-of-1), this work is presented as practice-based evaluation and does not constitute clinical diagnosis or treatment guidance.

4. Case Implementation (N-of-1) and Results

Across 14 available logs over the observation period, triage status was predominantly GREEN, with a short AMBER interval on 2026-03-10 to 2026-03-11, followed by return to GREEN status. No RED status was recorded.

Self-reported pain remained low overall, with reduced variability in later logs (early pain_worst up to 3–4; later generally 1–3). Functional tolerance showed improvement in standing and walking domains over the recorded period, while sitting tolerance remained context-dependent and lower without support. Sensitivity persisted in two recurrent contexts: morning flexion (e.g., getting up) and unsupported sitting.

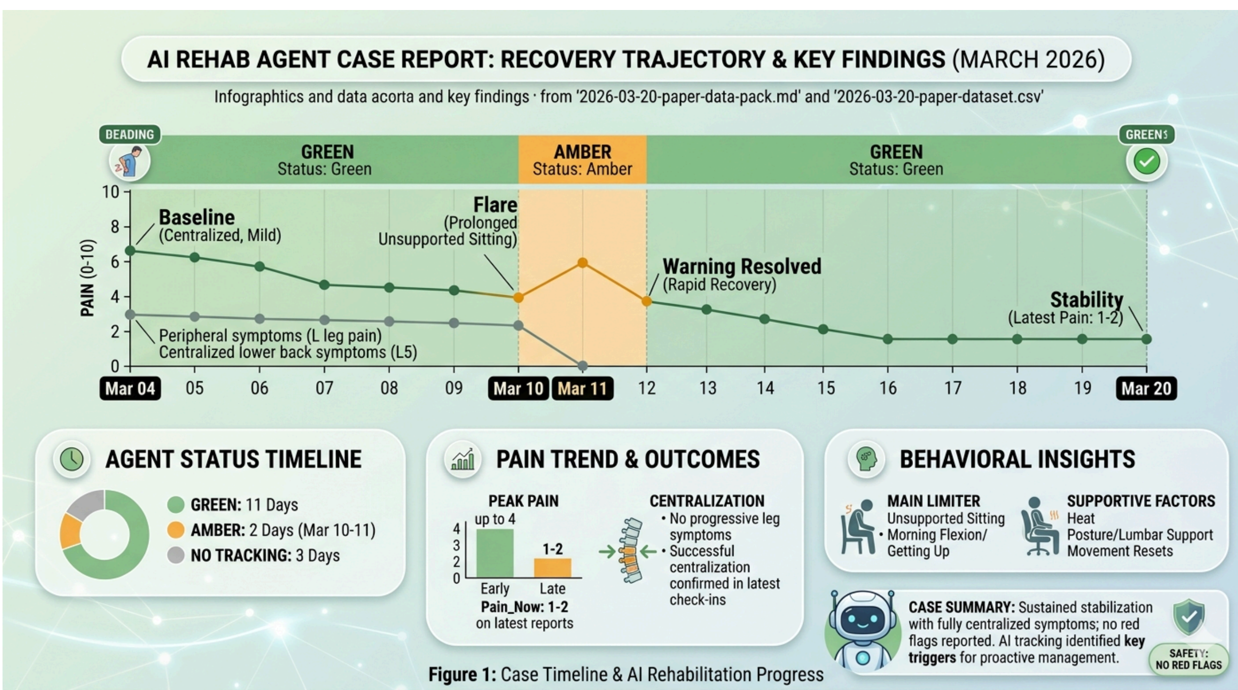

Figure 2: Daily status and outcome summary across the observation period. The figure illustrates the distribution of triage states (GREEN/AMBER), alongside trends in self-reported pain and functional tolerance.

No progressive leg symptom pattern or red-flag reports were documented in the available check-ins (Fig. 2).

5. Key Observations

Three main observations emerged from this single-case implementation:

First, recovery became more predictable as variability in reported symptoms decreased across days. Second, functional improvements in standing and walking were observed alongside persistent sensitivity in sitting and flexion, allowing for more targeted management rather than global activity restriction. Third, the structured, rule-based approach appeared to reduce day-to-day variability in decision-making, supporting more stable progression despite short-term symptom fluctuations.

6. Discussion and Future Integration

The primary contribution of this work is structured, safety-constrained daily decision support under uncertainty, rather than diagnosis or autonomous treatment.

In this single-case implementation, the framework provided a consistent approach to managing day-to-day variability and supporting progression decisions between formal clinical interactions. The emphasis on predefined safety rules and conservative progression aimed to reduce overreaction to short-term symptom fluctuations.

While this implementation focused on daily rehabilitation decision support, future work may explore integration with biomechanical modeling tools (e.g., OpenSim) and AI-assisted coding environments to link model-based insights with real-world functional data. This integration was not directly implemented in the present case.

7. Limitations

This work represents a single-case (N-of-1) implementation with self-reported outcomes. The framework is non-diagnostic and depends on the accuracy and consistency of user inputs. External validation and controlled studies are required to assess generalizability.

Findings may also be influenced by self-monitoring and expectation effects, and no external clinical adjudication of symptoms or classifications was performed.

Additionally, the system is intentionally conservative, which may limit progression speed in some cases.

8. Conclusion

In this N-of-1 implementation, a lightweight, safety-constrained, AI-assisted framework was associated with improved consistency and confidence in daily rehabilitation decision-making and may serve as a practical complement to clinical and biomechanical workflows.

Acknowledgment

The author acknowledges the use of AI tools in the development and preparation of this manuscript, including systems from OpenAI, Google (Gemini), and OpenClaw. These tools assisted with structuring, drafting, and refining the content. All final interpretations, analyses, and conclusions remain the responsibility of the author.

Keywords

AI-assisted rehabilitation, decision support, N-of-1, digital health, low back pain, self-management, agentic systems