

Flexor

Radison Akerman
Caffeine Hackathon
July 15th, 2025

Overview & Core Concept

Flexor is a web-based fitness app that generates personalized stretching routines using real-time soreness tracking, recent activity data, and intelligent muscle targeting. It's built for athletes, desk workers, and anyone who wants to move better without following rigid, one-size-fits-all plans. Users log in securely, land in a refined dark UI, and immediately gain access to an AI Routine Generator that delivers a dynamic, custom-duration mobility sequences crafted around their body's current needs.

Instead of navigating static exercise libraries, users interact with muscle zone sliders to indicate soreness or fatigue. The app then draws from a curated and extensible library of stretches, selecting movements based on their recovery goals, recent Strava activities, and past routines. The system avoids repetition, prioritizes relevant areas, and allows for deep customization.

Flexor's minimalist three-page structure — Routine, Body, and Log — makes for an experience that feels intentional and fast. Each stretch is visually clean, tappable, and adjustable. Logs show your recent rides, runs, strength sessions, and yoga flows, paired with the muscle groups they engaged. The result is an application that feels like a mobility coach that understands your body better with each session.

Purpose & Motivation

Most mobility apps on the market force users into linear plans. If you have hip tightness, you're prescribed a 30-day video series. But what if you're short on time today? Or sore in a different place tomorrow? Apps shouldn't force you to adapt to them. They should adapt to you.

As a triathlete and software engineer, I experienced this gap firsthand. Sitting long hours at a desk often led to anterior pelvic tilt and tight hip flexors. Biking left my calves and glutes stiff. Yet existing apps rarely considered these patterns, or made it easy to adjust routines on the fly. I wanted to build a tool that respected the fluid nature of real life and training, while also meeting a high bar for design and technical polish.

Problem & Response

Most recovery tools in the fitness space are overly rigid. They follow fixed timelines, assume uniform needs, and rarely acknowledge the variability of real-life training or sedentary patterns. Whether someone completes a high-intensity workout or simply spends the day seated at a desk, the experience remains unchanged. This lack of responsiveness makes recovery feel generic, uninspiring, and ultimately easy to skip.

Flexor takes a different approach by centering the experience on immediate context. Instead of guiding users through preset programs, it begins by asking how they feel and what they've done recently. By combining soreness inputs, time constraints, and logged activities, it creates short, meaningful routines that reflect the body's current state. Each stretch is selected for purpose, resulting in sessions that feel tailored rather than templated.

This adaptability encourages consistency. When a tool acknowledges effort, respects time, and adjusts to shifting priorities, users are far more likely to return. Flexor becomes not just a recovery app, but a trusted companion that evolves with the user. It invites participation by making recovery efficient, relevant, and intuitive.

Architecture & Approach

Flexor was developed using [Caffeine](#), a natural language-based platform that allows developers to build interfaces and logic through structured conversational input. Rather than editing source code directly, all changes were implemented by describing design and behavior goals in detail. This constraint demanded clarity in every decision and led to a system that is both intentional and highly composable.

At the core is a schema-driven model connecting stretches, muscle zones, activities, and user states. Each stretch includes metadata such as targeted regions, movement type, and estimated duration. Logged activities carry time, tags, and muscle impact values, which inform the AI routine engine. This structure enables Flexor to maintain a growing history of user input and adjust future recommendations based on both recent and cumulative behavior.

Routine generation follows a rule-guided, context-sensitive flow. The engine filters stretches by their relevance to current soreness, deprioritizes overused movements, and balances coverage across muscle groups. If a user reports tight calves after a ride, for example, Flexor will surface posterior chain stretches but introduce enough variation to avoid repetition. This logic maintains freshness while keeping the routines focused and efficient.

The interface was assembled through iterative refinement, with layout, spacing, transitions, and component structure described through precise natural language prompts. This process shaped a lightweight, visually calm environment that reduces distraction and encourages flow. Small touches, such as animated sliders or seamless stretch reordering, reflect a careful attention to microinteractions. Although the underlying architecture is modular and extensible, the user experience remains cohesive and focused.

Audience & Use Cases

Flexor targets anyone who wants to improve mobility in a smarter way. While the broader market includes general users, the core audience is:

- Athletes and fitness enthusiasts
- Desk workers with posture or flexibility issues
- Users already tracking workouts via Strava, Garmin, or Apple Health
- People rehabbing injuries or following physical therapy routines

Flexor gives these users a powerful way to stretch better without needing to think too hard. It lowers the barrier to recovery by meeting people exactly where they are.

Roadmap & Potential

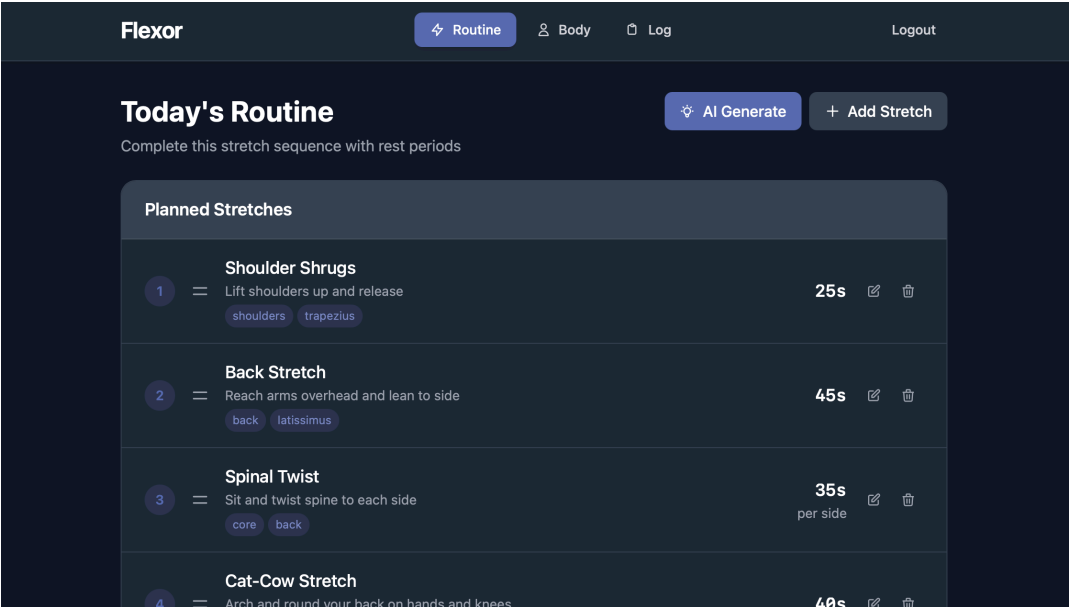
With continued development, Flexor could integrate tightly with services like Strava or Apple Health to auto-import activities and infer soreness zones. Future features might include:

- Voice input: "My hamstrings are wrecked" automatically updates sliders
- Short video demos for every stretch
- Real-time notifications: "Stretch your back in 5 minutes between meetings"
- Progress analytics over time: "You've reduced hip soreness 30% this month"
- Community-created stretch packs based on sport, job, or routine

Flexor is designed with scale in mind. Its modular architecture supports rapid feature expansion and data integration. The stretch library is dynamic and user-expandable. With strong user engagement, Flexor could even power API-driven recovery recommendations for wearables or performance coaches.

Interface

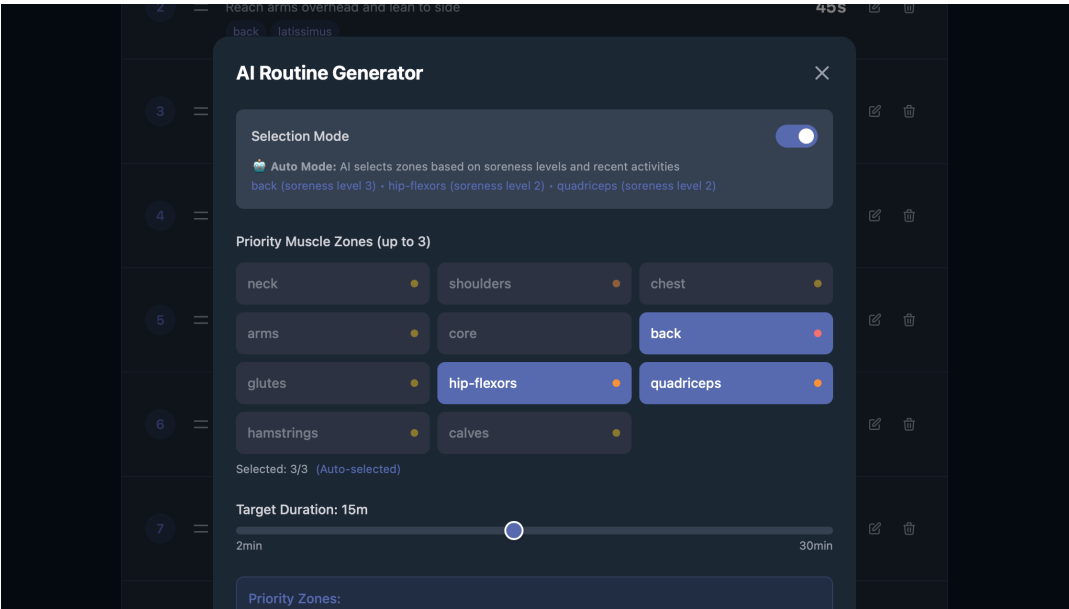
This section presents a glimpse into Flexor’s preliminary visual experience and design language.



Routine Page

The central hub for composing and reviewing stretch sequences. Users can drag to reorder movements, tap for details, or generate a new routine based on current needs.

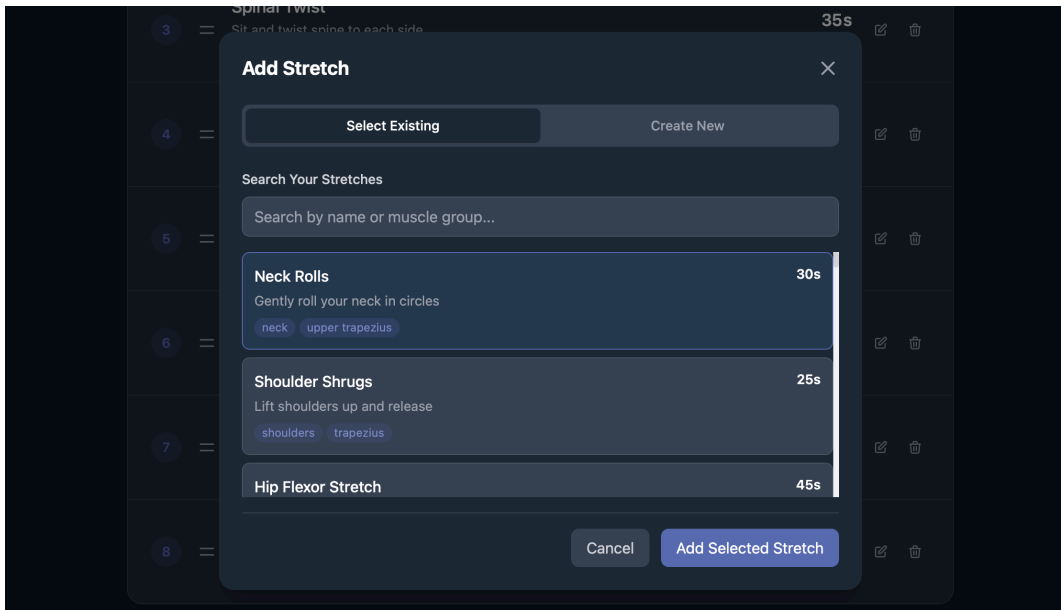
...



Routine Generator Dialog

A focused modal for customizing routine generation. Users select desired duration and muscle zones to receive a tailored, soreness-aware routine.

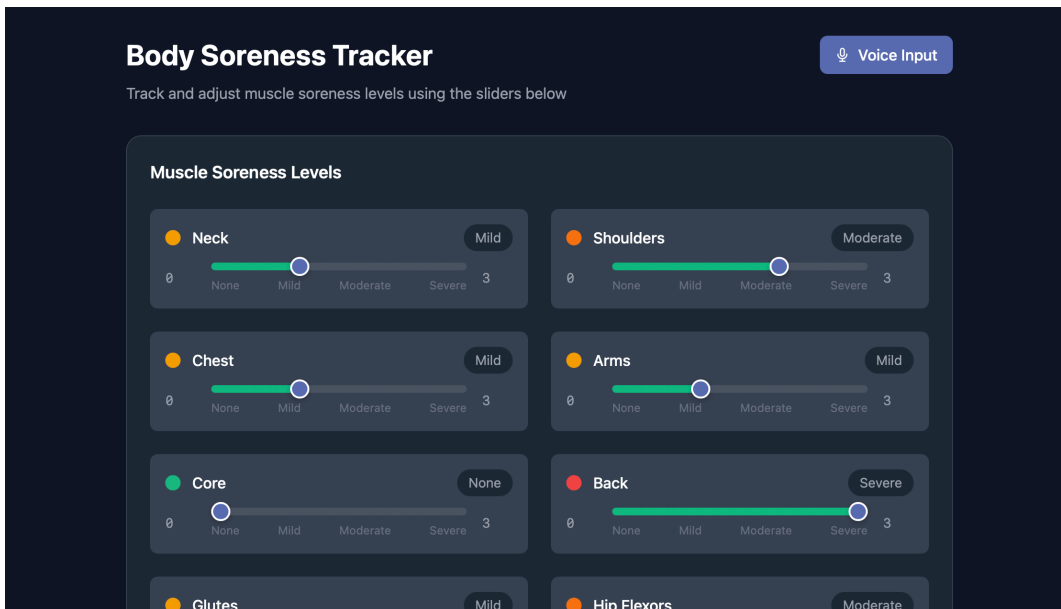
...



Manual Add Stretch Dialog

Allows users to add stretches manually from the library. Each entry includes metadata such as target muscle groups and duration, offering full control over the sequence.

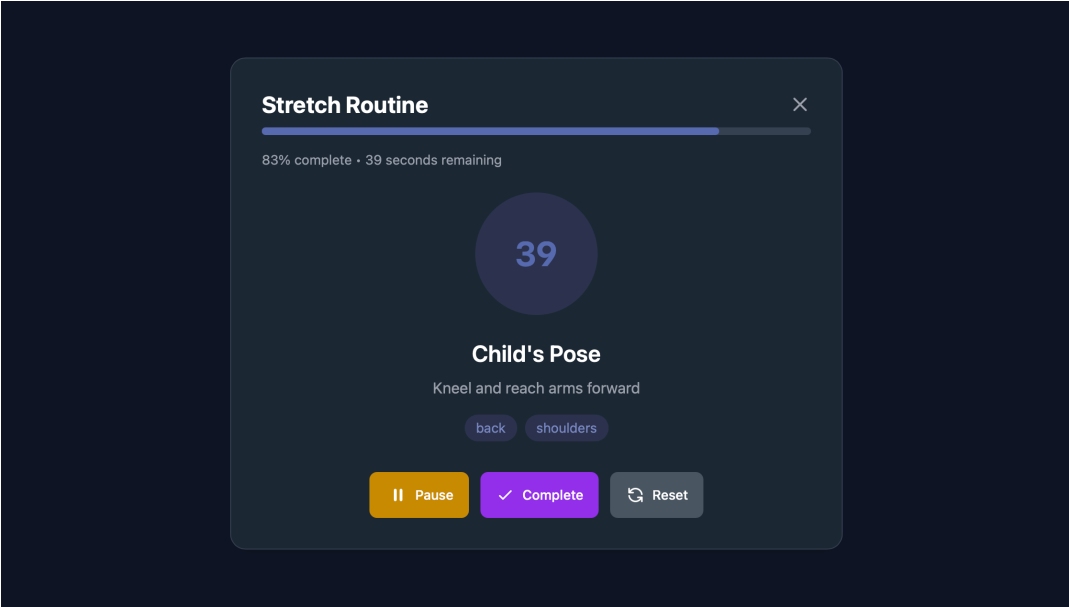
...



Body Page

A clean, slider-based interface for updating soreness across muscle zones. Inputs on this page directly influence routine recommendations and stretch targeting.

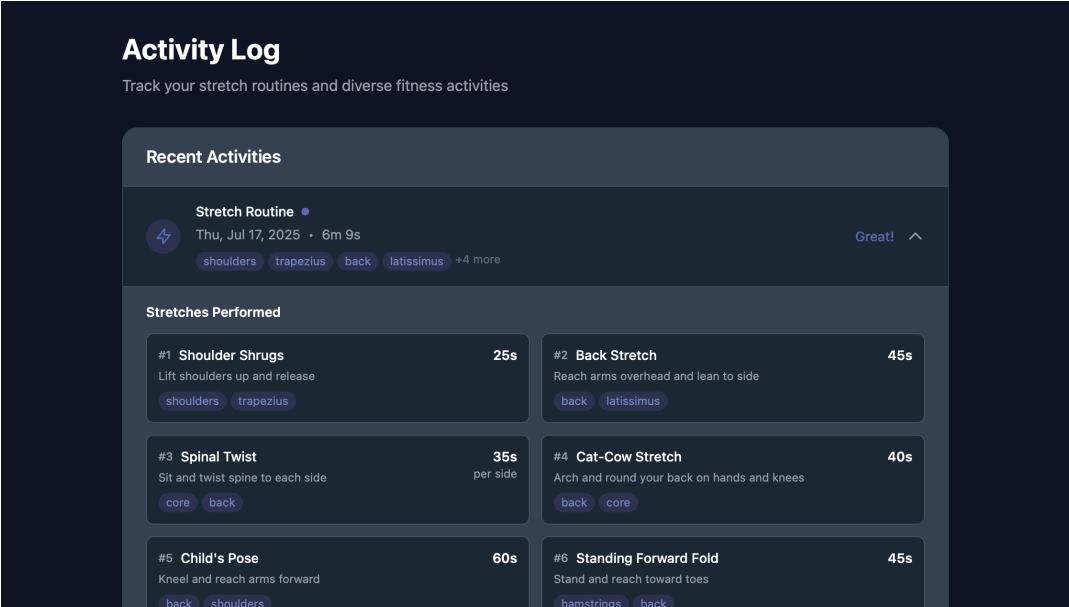
...



Stretch Routine Walkthrough

A guided interface that steps through each stretch with clear visual emphasis. Ideal for users following along in real time with minimal distraction.

...



Activity Log Page

Chronicles recent activities, such as rides, runs, and strength sessions, tagged by muscle engagement. This log feeds the AI engine, keeping future routines context-aware and varied.

...