siliconmint

Scaling Mental Health Access with AI-Powered Clinical Workflows

Improving mental health care in underserved areas using Google Gemini and precision prompt engineering

We partnered with a U.S.-based behavioral health platform whose mission is to make high-quality mental health care more accessible — especially for patients in rural and underserved areas. By integrating Google Gemini's AI capabilities and developing a high-precision prompt engineering framework, we enabled faster, more consistent clinical decision-making at scale. The result: more patients served per day, less administrative burden, and broader access to care.

siliconmint

Business Challenge

The client’s vision was to bring structured, AI-supported mental health care to communities that typically lack psychiatric infrastructure.

To succeed, they needed to:

- Drastically improve the efficiency of intake and diagnostic workflows across clinics
- Support clinics operating in low-resource or geographically remote environments

The challenge was twofold:

- Enable intelligent decision-support without overloading clinical staff
- Accelerate clinical operations while maintaining care quality

siliconmint

Technical Challenge

Working with AI in a clinical setting required extreme attention to accuracy, safety, and contextual nuance.

Core challenges included:

- Building smart, adaptable prompts to engage Google Gemini with clinically relevant questions
- Supporting 50+ assessment formats, including complex multi-section tools and free-text responses
- Integrating patient-recorded videos into the clinical workflow via speech-to-text and structured extraction
- Preventing vague or risky suggestions

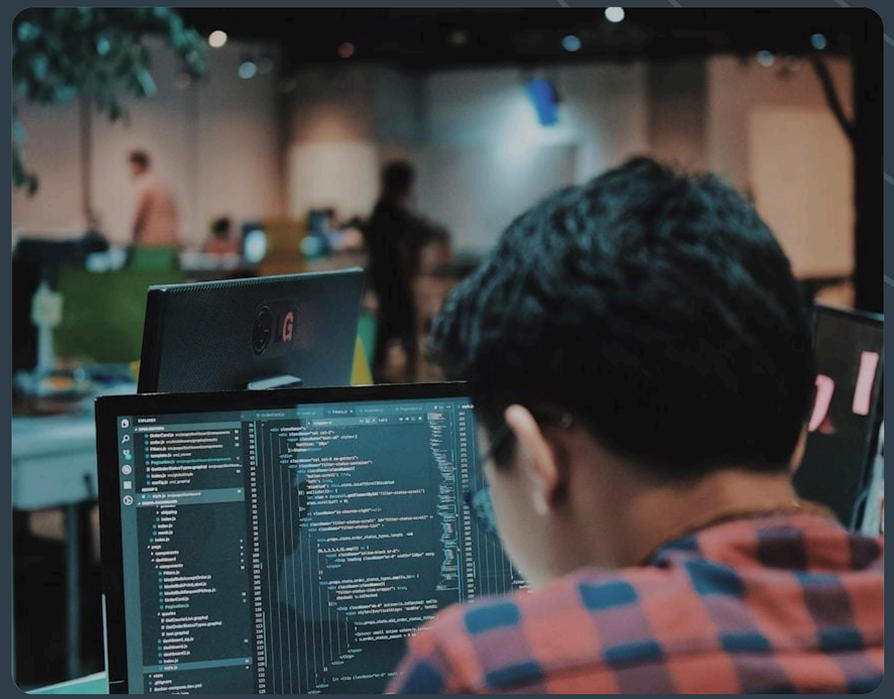

siliconmint

Model definitions

MODEL NAME	ARCHITECTURE	PROVIDER	PURPOSE	INPUT	OUTPUT	USAGE CONTEXT
MedSuggest-XGB	XGBoost (Gradient Boosting Trees)	Gemini ML Platform	Suggests probable diagnoses based on structured patient data	EHR vector: symptoms, vitals, history (~100 features)	Top-N diagnosis candidates with confidence scores	Triggered during in-person or async consultation
SafeCheck-Net	Shallow Neural Network + Rule Engine	Custom (Python, ECS)	Validates doctor-entered diagnosis for contradictions	Diagnosis input + medication and condition context	Risk flags, validation messages	Real-time validation in physician UI
VideoDiag-Gemini	Multimodal Transformer + Speech2Text Pipeline	Gemini Multimodal API	Analyzes patient-recorded video responses and suggests possible conditions	Short recorded videos (60–120 sec), transcribed text	Suggested diagnosis + confidence score + structured symptom record	Used for remote patients or symptom triage intake; adds to clinical record

Runtime characterictics

MODEL NAME	AVG INFERENCE TIME	RETRAINING FREQUENCY	CONTEXT VOLUME	EXPLAINABILITY	REAL-TIME COMPATIBLE
MedSuggest-XGB	~250 ms	Weekly	Structured data vector (~100 vars)	✅ Feature importance (SHAP)	✅ Yes
SafeCheck-Net	~300–400 ms	Quarterly (rules), Rare retrain	~20–30 input fields	✅ Rules + traceable logic	✅ Yes
VideoDiag-Gemini	~2–3 sec (incl. ASR)	No retraining (Gemini prompt)	~1–2 min spoken input + transcript (~300 tokens)	✅ Natural-language explanation (prompt)	⚠️ Yes, but post-recording, not fully live

Tech Stack

React, GraphQL, Node.js,
real-time assessment parser, clinical validation module

Node.js

React

GraphQL

siliconmint

Our Solution

We engineered a dynamic AI interaction layer tailored for behavioral health assessments, enabling real-time communication with Google Gemini. Our prompt engine was built to accommodate a wide range of assessment types, patient inputs, and clinic-specific workflows.

Key Features

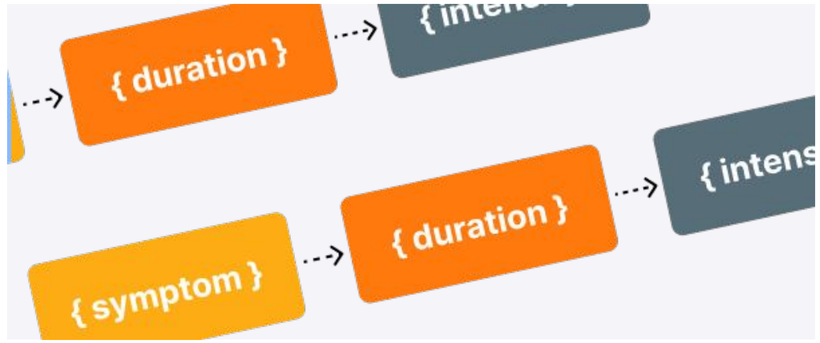

Context-aware prompt templates:
Tailored prompts for clinical accuracy

Our dynamic prompts are adapted to each assessment type, patient input style, and care environment — ensuring AI responses are precise, clinically meaningful, and easy to interpret.

AI Risk Detection:
Filtering vague or unsafe responses

Every AI-generated output is validated in real-time to catch ambiguity or clinical risk — preventing low-quality suggestions.

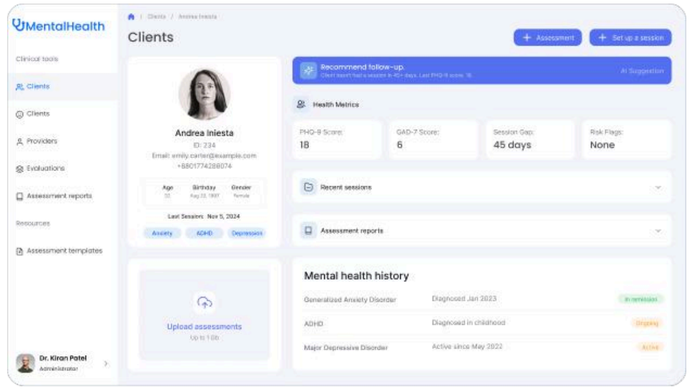

Real-Time Decision Support:
Insights where they matter most

Clinically relevant AI suggestions appear directly in the interface — no context-switching, no delays.

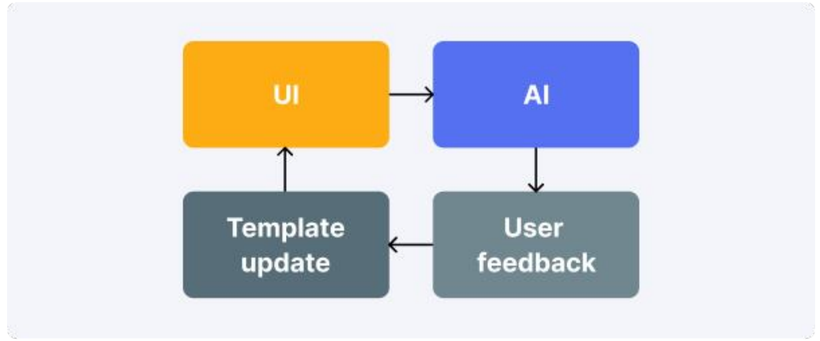

Continuous Prompt Optimization:
Learning from real-world use

We monitor how clinicians interact with prompts and AI output, continuously refining templates to improve clarity, efficiency, and relevance.

Seamless System Integration:
Compliance-ready and workflow-friendly

Our AI layer integrates fully with existing clinical systems, ensuring smooth adoption, data consistency, and regulatory compliance.

This allowed providers to spend less time on documentation and interpretation, and more time with patients — especially important in rural clinics where staff are limited and every minute counts.

siliconmint

Clinically-Driven Product Development

The product team maintained an active feedback loop with clinicians and operational staff at partner clinics. Product Owners conducted regular interviews with healthcare providers to understand pain points, unmet needs, and opportunities for improvement. Many new features — from UI enhancements to AI output tuning — were directly driven by this ongoing, clinician-informed discovery process.

We demo, we listen, we improve.

1. **Up to 10 customer feature requests collected weekly**
   Direct input from clinicians ensures we build what truly matters.

2. **Requests prioritized by frequency and urgency**
   The most common pain points rise to the top of our backlog.

3. **Feasibility and impact assessed before each sprint**
   We commit only to features that are timely, useful, and buildable.

4. **Features developed and tested in tight cycles**
   Delivery is fast, focused, and aligned with clinical workflows.

5. **Live demos conducted with real providers**
   We validate value in real time and gather actionable feedback.

6. **Iterative delivery with post-launch feedback**
   We track adoption and adjust based on how features perform.

siliconmint

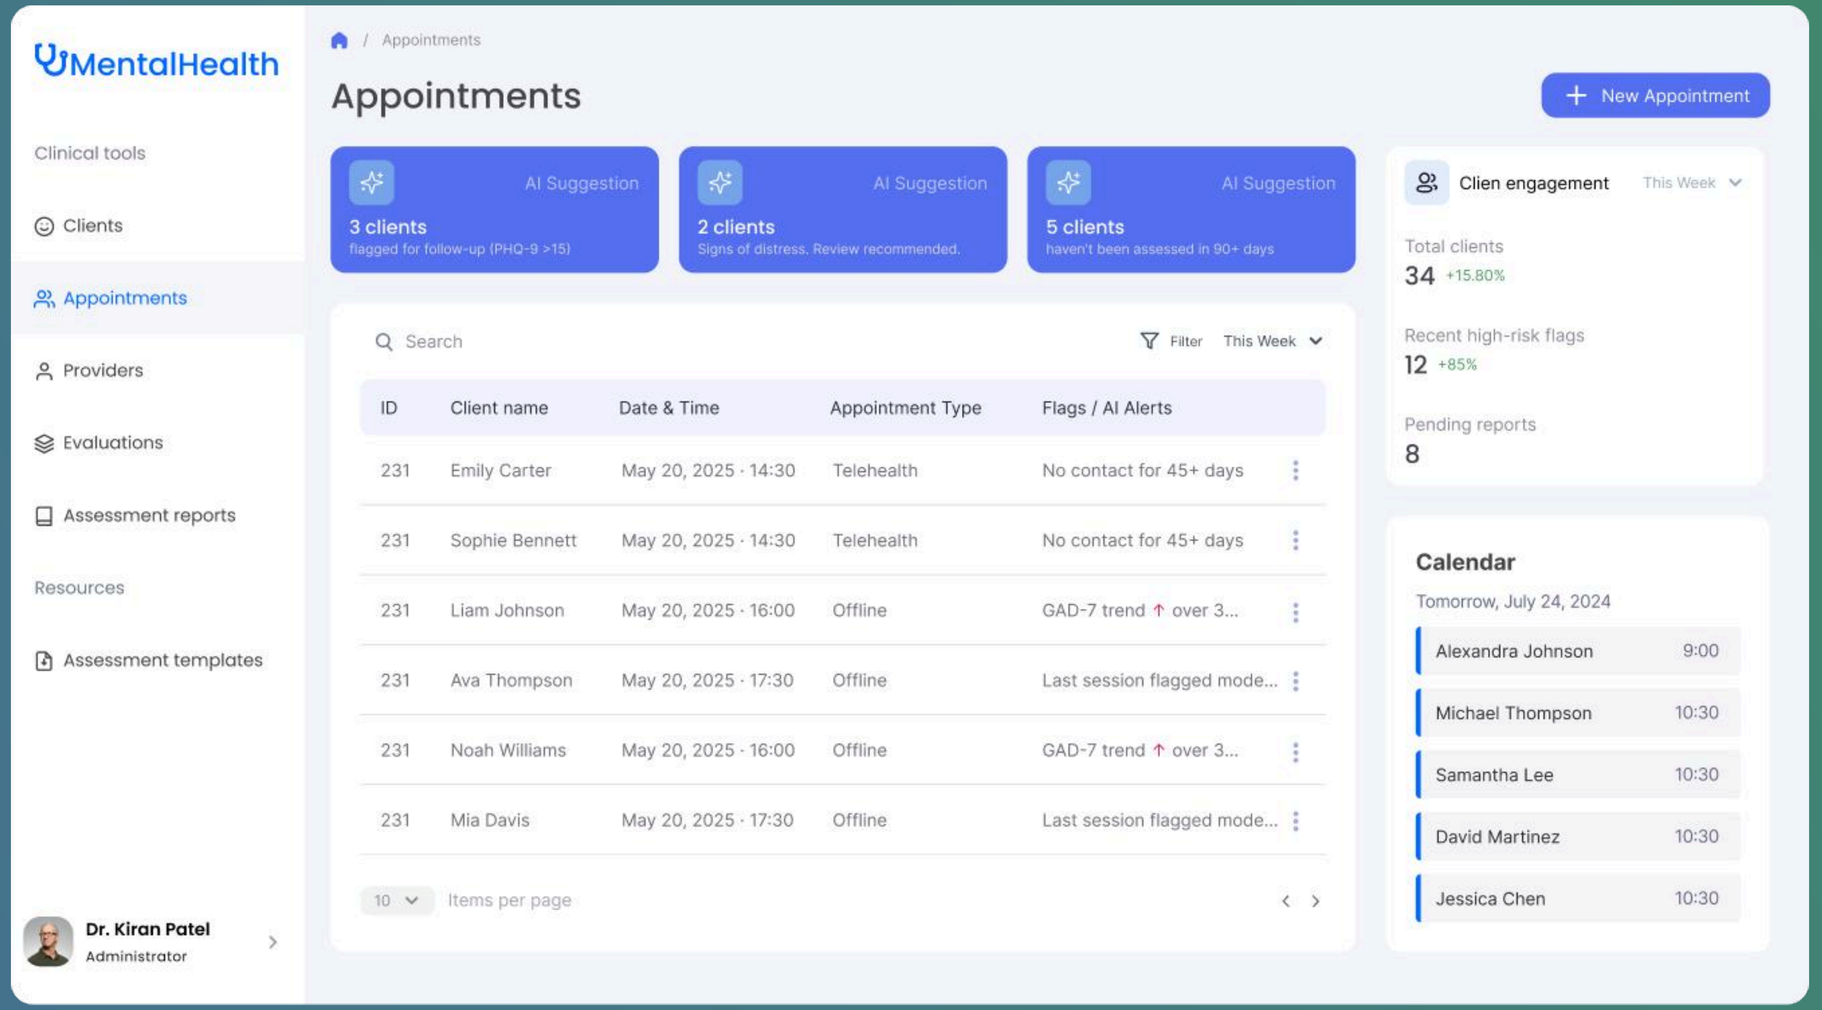

siliconmint

Impact

Expanded access to care

in rural areas through increased throughput

3–4× faster clinic onboarding

enabling rapid scale

~50% reduction

in time to interpret assessments

>90% provider validation accuracy

on AI diagnostic suggestions

50+ assessment types supported

enabling consistency across providers