

Junior — UI/UX Research Report & Design Direction

Version 1.0 · August 3, 2026 · Companion to: Junior Screen Inventory, Competitive UX Teardown (HTML), and hi-fi mockups

1. Executive Summary

This report answers three questions: (1) where AI-native product UX is heading and what that means for a healthcare forms platform; (2) what the competition actually looks like on screen; and (3) the recommended design direction — a **chat-first shell** with **voice as a first-class input**, wrapped in a new "Calm Clinical + AI Glow" design language.

The headline recommendation. Build Junior so the AI conversation is the center of the product: one persistent omnibox (typed or spoken) that is simultaneously navigation, command search, and prompt; workspaces (Library, Designer, Governance, Analytics) that open as canvases around the conversation; and AI output rendered as interactive cards from a fixed component catalog. Then draw one hard line: **every regulated moment — signature, consent, publish, payment, permissions — drops into deterministic, fixed-layout UI. Chat drafts; humans sign.** This combination is both the most differentiated position in the market and the only chat-first architecture a hospital compliance officer can approve.

The market opening. No competitor — Epic, Meditech, Phreesia, Taylor, Interlace — has applied AI to forms themselves. Their AI investments target phones (Phreesia VoiceAI), ambient scribing (Epic, Meditech/Nuance), and assistants (Epic's Emmie/Art/Penny). AI form generation, PDF-to-form conversion, adaptive intake, and AI-managed language variants are unclaimed. Junior's \$15 AI subsystem, made visible through this design, is the differentiator.

Three corrections to the thesis documents (from the competitive research):

1. **Phreesia acquired Access eForms (Aug 2023).** "Access eForms Passport" is now a Phreesia product and is the MaaS-compatible forms solution in Meditech's marketplace. The vendor the thesis credits with "winning the market inside Meditech MaaS deals" belongs to the biggest intake incumbent.
2. **Interlace Health is the rebrand of FormFast, not Access** — and its forms products lost MaaS compatibility. Meditech is now bundling a direct competitor's (Phreesia's) product: a structural opening for an independent, white-label, cloud-native forms layer.
3. **Phreesia is giving its intake platform away free through Dec 31, 2026** (paid from Jan 2027). Expect price-zero competition in ambulatory; Junior's wedge is acute care, partners, and white-label.

2. Where AI-Native UX Is Heading (2025-2026)

Eighteen patterns were identified across ChatGPT (Canvas, Apps SDK, unified voice), Claude (Artifacts, Cowork), Microsoft Copilot (Pages, BizChat, Copilot OS concept), Gemini in Workspace, Perplexity, Notion AI, Vercel's generative-UI stack, and the clinical-AI vendors (Abridge, Nuance DAX, Suki). The ones that matter most for Junior:

2.1 The shell patterns

- **Side-by-side canvas** (ChatGPT Canvas, Claude Artifacts): conversation in one pane, the durable work product in an adjacent editable surface. The chat is ephemeral and directive; the artifact is stateful. Junior's Form Designer is exactly this: chat left, form canvas right.
- **Apps/cards in conversation** (OpenAI Apps SDK): interactive components render inline in the thread — expand in place rather than navigating away. Junior renders form previews, analytics tiles, and approval cards inline; each has an "open as workspace" affordance.
- **Chat inside the app — the inverse** (Gemini side panel, Copilot in M365): when a user is deep in a deterministic screen, the conversation stays docked with full context of what's on screen. Junior's workspaces all carry a docked chat.
- **Promote chat output to a governed object** (Copilot Pages): any AI draft gets one-click "save to library" — converting an ephemeral chat result into a versioned, permissioned, governed object. Chat scrollbar is never the system of record.
- **Command palette fused with AI** (Notion, Linear): one Cmd-K omnibox that accepts fuzzy command search and free prose — solving chat's discoverability problem by listing what the system can do.

2.2 Generative UI, safely

The engineering pattern behind inline cards: an LLM tool call maps to a typed component with three render states (loading shimmer → rendered → error). The critical constraint from practice: **the AI chooses and populates components from a curated catalog — it never free-renders novel UI**. Deterministic components, probabilistic selection. This is precisely the safety posture healthcare needs, and it makes the AI's output testable.

2.3 Voice

- **Voice is a modality of the same surface, never a separate room.** OpenAI collapsed ChatGPT's separate voice screen into the main chat: speech in → streaming transcript + cards out. Junior does the same — no modal voice "orb."
- **Push-to-talk / tap-to-toggle is the default.** Clinical environments are noisy and privacy-sensitive; open-mic modes get interrupted and erode trust (OpenAI's community demanded push-to-talk). Wake word and ambient capture are explicit opt-ins with persistent recording indicators.
- **Show what you heard.** Live waveform + real-time transcript while listening; voice-filled fields carry a visible "from voice" marker.

- **Voice fills drafts; it never submits.** Speech populates fields and stages actions; anything consequential is confirmed by tap/click. Suki-style voice commands ("next section," "read back allergies") layer on top of dictation.
- **v1 constraint honored:** §15.1 forbids AI paths that could take PHI. Ship the voice affordance visible but policy-disabled on PHI surfaces until a BAA-compliant AI configuration exists; enable first on non-PHI surfaces (designer, library search, analytics).

2.4 Trust & audit affordances (the healthcare-grade layer)

- **Linked evidence** (Abridge's pattern): every AI-populated field, mapping, or claim links to its source — the document passage, the model property, the utterance. Unsourced content is explicitly badged, never silently blended.
- **Draft/verified state machine:** AI output is born **draft** (gradient border + ✦ badge); only a named human transitions it to verified/published — and the glow disappears permanently on approval. Verification is visible state, not a toast.
- **Approval inbox as a first-class surface:** all human-in-the-loop gates aggregate into a queue with previews, diffs, and role-based routing — approvals are workflow objects, not lost chat messages. This maps 1:1 onto Junior's Forms Governance.
- **Autonomy dial:** per-action-class autonomy configuration (auto / needs-approval / blocked), mirroring the three-tier model already written into the Forms Library and Marketing engine PRDs. Compliance officers set, e.g., "copy edits: auto; anything touching consent language: manual forever."
- **Binary confidence, not percentages:** low-confidence AI-filled fields get an amber "verify" treatment and are excluded from bulk-accept.

2.5 When chat-first fails — and the hybrid answer

The strongest research critique ("The Conversation Trap," NN/g): conversational input is slow for known tasks, loads working memory, and hides capabilities. The gate questions: *could this be done faster by clicking?* and *does the value lie in the outcome or in the visible human process?* Where process carries legal meaning — signatures, consent, publishing — conversation is the wrong proxy. Hence Junior's three-layer rule:

4. **L1 Shell** — chat/omnibox everywhere, for intent capture, search, drafting, Q&A.
5. **L2 Workspaces** — direct-manipulation screens for known, repeated tasks (the designer's drag-drop, the governance queue), each with docked chat.
6. **L3 Deterministic surfaces** — fixed-layout ceremonies for regulated moments: signature adopt/apply, PIN entry, EULA, payment, publish confirm, permission changes. Identical every time; trainable; auditable; never generated.

Every AI capability must also exist as a visible button or menu (scoped actions). Chat is an accelerator, never the only path.

3. Competitive Summary

Full visual teardown with screenshots is in the companion HTML document. The essentials:

Vendor	UX model	AI story	Weakness Junior attacks
Phreesia	No-app SMS links; stepper intake; schedule-centric staff hub	VoiceAI phone agent (2025)	Form changes require vendor services; long flows frustrate elderly patients
Access eForms (Phreesia)	Meditech-launched web portal; HL7 prefill; barcode archive	None of its own	Dated UI; not cloud-enabled; brand absorbed
Interlace Health (ex-FormFast)	Forms on Demand library; print + barcode; downtime niche	None ("consent for others' AI")	"Antiquated" builder; lost MaaS forms slot
Epic	MyChart eCheck-in stepper; Welcome kiosks; analyst-built questionnaires	Emmie/Art/Penny, ambient charting	New form = certified-analyst IT ticket
Meditech	Expanse tiles; Navigator AI pane; partner-delegated forms	Navigator, ambient docs (Google)	Weak patient UX (2.9★ portal); no native builder
Taylor Healthcare	Procedure-picked consent assembly; signature pads; print DNA	None	Consent niche; no builder; opaque product
Jotform	20k-template catalog; Build/Settings/Publish ; AI Agents on 14+ channels incl. phone	Strongest in category	No EHR integration, no healthcare depth
Typeform	One-question fill craft; motion as wayfinding; Clarify-with-AI	AI Engagement Platform (2025)	Not a healthcare/compliance product

Vendor	UX model	AI story	Weakness Junior attacks
DocuSign	Guided next-field ribbon; adopt-once signature; Focused View embedding	Iris: signer summaries + Q&A, AI Web Forms	Documents, not data; no EHR writeback

Top patterns to steal: (1) form-as-agent duality — one form model, three fill surfaces (page / card / conversation); (2) prompt→form lands in the normal editor, never a separate AI mode; (3) skip-to-next-required ribbon with a "review entire form" escape hatch; (4) adopt-once signature ceremony; (5) Focused View chrome-less embedding for white-label; (6) in-form AI answering ("why do you need my SSN?"); (7) dual-taxonomy template gallery with live preview; (8) Typeform's fill craft (big type, keyboard, motion, piping); (9) celebrate prefill ("we filled 14 fields from your record"); (10) schedule-centric staff command center.

4. The New Design Language — "Calm Clinical + AI Glow"

Design mantra from the PRD: *Safe, Effective, Compliant, Feature Rich... yet Beautiful and Easy to Use*. The language expresses it in four rules.

4.1 Foundations

- **Surfaces:** warm paper white (#FAFAF7) base; pure-white cards; soft sage-tinted wells (#F3F4F0). Light theme default (healthcare users, print-adjacent work); dark theme later.
- **Ink:** near-black #171A1E; secondary #4A5058; tertiary #878E97. Body 14–15px at 1.7 line height.
- **Brand primary — "Marine":** a calm clinical teal (#0E6E65 for actions, #DDF0ED tints). Reads trustworthy and healthcare-native without being hospital-beige.
- **Type:** Sora (display — geometric, modern, light weights for greetings and numbers) + Inter (UI) + JetBrains Mono (code/expressions).
- **Shape & depth:** 12–16px radii; layered soft shadows; no hard borders on primary surfaces.
- **Motion:** 100–300ms state transitions only — shimmer for retrieval, streaming for generation, pulse for agent execution, auto-scroll as wayfinding. No decorative animation.

4.2 The AI provenance rule (the signature move)

A single vivid gradient — indigo #6D5EF3 → cyan #19B8D9 — and the ✦ glyph are reserved **exclusively** for AI-originated content and actions. AI drafts carry a gradient border and ✦ badge; the moment a human approves or edits, the glow is removed permanently. Human-entered content is always visually plain. Ordinary chrome never uses the gradient. Result:

provenance is legible at a glance across the entire product — a compliance feature expressed as a visual identity.

4.3 Status & semantic color

Form lifecycle chips: Not Started (gray) → Prepped (blue) → In Progress (amber) → Complete (green). Semantic: success #1E8E5A, warning #B7791F, danger #C43D3D, info #2B6CB0 — each with a tint. All pairs validated for WCAG AA on their backgrounds.

4.4 White-label strategy

The design system is token-driven: partners override the Marine primary, logo, and radius scale; the AI gradient and provenance rules remain constant platform-wide (they are trust infrastructure, not brand). RTL mirroring is built into the layout primitives (the Viewer already flips for Arabic).

5. The Chat-First Shell — How Junior Is Structured

- **Left rail:** Approvals & Tasks inbox (badge counts), AI-drafts-to-review, pinned workspaces, recent threads, identity (SSO via EHR).
- **Center:** the conversation. Greeting + glance cards for the day; messages; AI responses composed of catalog cards (tasks, charts, form previews, diffs). Bottom composer = the omnibox: type, ⌘K, or hold-to-talk.
- **Right / canvas:** workspaces open here (Library, Designer, Governance, Analytics) — full direct-manipulation screens, each with docked chat. Small cards expand in place; big work promotes to canvas.
- **Context transparency:** the composer always states what Junior can see ("your Forms Library, governance queue & analytics — no PHI in chat").
- **Persona defaults:** Designers land in design-centric suggestions; Admins see billing/usage; Governance approvers see their queue; Partners see tenant roll-ups. Same shell, different first cards.

6. Mockups Delivered

7. **01-portal-home-shell.html** — the chat-first home: rail, glance cards, voice-transcribed user message, AI reply with a deterministic task card + an AI-drafted analytics card (with linked-evidence footer), suggestion chips, omnibox with mic.
8. **02-forms-library.html** — Library workspace: AI natural-language search ("find all consents missing the Chinese variant"), AI answer strip with bulk-draft action, filter rail, category list, form cards with language coverage chips and governance status, docked contextual chat.
9. **03-form-designer.html** — the AI canvas: chat pane driving changes (plan card with staged steps, Accept/Refine/Discard), element toolbox with Patient-model drag chips, form canvas with prefill-tagged fields and an AI-drafted conditional section awaiting acceptance, properties panel with function binding (getAge · On First Load), language

tabs with out-of-sync indicators, device preview toggle, Request-publish (governance) action.

7. Recommendations & Next Steps

10. **Adopt the three-layer rule (L1/L2/L3) as an architectural principle** in the PRD — it resolves Q-15 (Help-bot scope: the bot IS the omnibox, scoped actions are its buttons) and de-risks compliance review of the AI subsystem.
11. **Add "AI provenance" to the design-system spec** as a hard rule; it operationalizes the PRD's audit posture visually.
12. **Voice rollout order:** designer/library/analytics first (no PHI), viewer form-fill later behind the BAA-compliant AI configuration (§15.1).
13. **Update the thesis docs** for the Phreesia/Access/Interlace corrections and the free-through-2026 pricing dynamic; sharpen the wedge language toward acute care + white-label partners.
14. **Screen inventory as build backlog:** the companion inventory enumerates 64 screens / ~280 states with popups and variations — structured to hand directly to AI build agents, one workspace at a time.
15. **Extend mockups next to:** Governance review (the deterministic HITL surface), the patient Viewer (Typeform-grade fill + DocuSign-grade signing), and the SoF embedded Forms tab.

Appendix A — Key Sources

AI-native UX: OpenAI (Canvas; Apps in ChatGPT; unified voice), Anthropic (Artifacts; Cowork), Microsoft (Copilot Pages/BizChat; Copilot OS concept; "Hey Copilot"), Google (Gemini Workspace side panel), Vercel (AI SDK generative UI; json-render), Notion 3.0, Perplexity citation model, NN/g "Scope in Generative AI," "The Conversation Trap" (designative.info), StackAI & LangGraph HITL patterns, shadcn.io AI components.

Clinical voice/AI: Abridge (Linked Evidence), Nuance DAX/Dragon Copilot, Suki (voice commands in Epic), Freed (audio deletion policy), Epic UGM 2025 coverage (Healthcare IT Today, CNBC).

Competitors: phreesia.com (+ pricing, VoiceAI & Access-acquisition press), interlacehealth.com, accessefm.com, ehr.meditech.com (Expanse, MaaS, vendor listings), taylor.com (iMedConsent, VA), mychart.org + Epic app-store listings, jotform.com (AI Agents, HIPAA, templates), typeform.com (AI Engagement Platform, Formless), docusign.com (signing UX, Focused View, Iris), fillout.com, feathery.io, formstack.com; G2/Capterra/GetApp/TrustRadius reviews; Becker's Hospital Review.