Stanford Online / Capstone Project
UI/UX Design for AI Products

Advisor Hub AI

1 / Problem Identification and Research
2 / Concept Development and User Interaction
3 / Ethics and Societal Impact Review

studio dwd®
Visual & Product Design / www.studiodwd.com / www.linkedin.com/in/davidwdieter/

1 / Problem Identification and Research/ Problem statement

In the life insurance business, advisors and case managers must collect and reconcile data from many disconnected sources—Salesforce (CRM), Gmail/Google Workspace, Google Drive, carrier portals and illustration systems, plus client medical and financial documents—to underwrite cases, design products, and prepare client proposals and reviews. The current process relies heavily on manual data gathering, rekeying, and spreadsheet or PowerPoint assembly, which is slow, error-prone, and hard to scale, especially when comparing multiple carriers and plan designs for a single client.

Regulatory and carrier compliance add another layer of work: for new sales and annual reviews, teams must pull policy data from multiple systems, run new and in-force illustrations, check against internal guidelines and regulations, and document recommendations in a compliant, client-friendly format. As volumes grow, this manual approach limits advisor capacity, delays client responses, and increases operational and compliance risk

Pain points

Fragmented data: Policy, medical, financial, and illustration data live in separate internal and external systems (CRM, email, cloud storage, carrier portals), requiring extensive manual search, download, and copy-paste work for each case.

Manual proposal creation: Advisors must run multiple carrier illustrations separately, then manually compare options and build a consolidated proposal or presentation for the client.

In-force and annual review burden: For reviews, staff repeat the same multi-system data gathering and illustration work, then rebuild Excel and PowerPoint reports from scratch.

Compliance friction: Checking recommendations and reports against regulatory rules and internal guidelines is largely manual, increasing the risk of missed disclosures, inconsistent documentation, and failed audits.

Inconsistent data and errors: Manual rekeying of values between systems (e.g., from carrier portals into Excel or CRM) introduces errors that can affect underwriting decisions and client trust.

Limited advisor capacity: High volumes of small, repetitive tasks (document chasing, formatting, follow-up emails) reduce time available for advice, planning, and client conversations.

AI opportunities (intelligence augmentation)

Unified data extraction and case workspace
Use AI document and data extraction to pull key fields from carrier PDFs, medical records, and financial statements directly into a structured case record in Salesforce or a case hub. Deploy connectors plus an AI "case assembler" that, given a client or policy ID, automatically gathers latest data from CRM, email threads, carrier portals, and cloud folders into a single workspace for the advisor to review and approve.

Assisted illustration and product design
Build an AI assistant that reads underwriting inputs and client goals, then recommends which carriers/products to illustrate, pre-populating illustration request templates and flagging missing information for the human underwriter to obtain. Use AI to ingest the resulting carrier illustrations and generate a normalized, side-by-side comparison of premiums, guarantees, riders, and assumptions, with editable narrative explanations for the advisor.

AI-driven proposal and presentation generator
Implement a proposal generator that converts multiple quotes/illustrations into a single client-ready PDF or slide deck, highlighting key differences, trade-offs, and suitability rationale while allowing the advisor to edit wording and emphasis. Enable "smart templates" that automatically pull the right charts, disclosures, and language blocks based on product type, client segment, and jurisdiction, again leaving final approval with the advisor or compliance.

Annual review copilot
Provide an AI "annual review prep" bot that, on a schedule, assembles current policy data, recent performance, new in-force illustrations, and relevant life events from CRM notes into a draft review packet for each client. The bot generates a concise advisor brief (what changed, risks, opportunities) and a client-friendly summary report, which the advisor can adjust before sending or presenting.

Embedded compliance and audit support
Use NLP-based compliance checking to scan proposals, emails, and review reports for required disclosures, prohibited phrases, and missing documentation, surfacing issues to advisors and compliance before anything goes to the client. Maintain an AI-generated, human-reviewable audit trail that links data sources, illustrations, recommendations, and client communications so that regulators and carriers can see how each recommendation was formed.

Workflow automation and smart reminders
Deploy AI agents to handle repetitive outreach and follow-ups: requesting missing medical/financial documents, reminding clients to sign forms, and nudging carriers on outstanding requirements, while routing complex responses to humans. Use predictive models to prioritize cases and reviews that are at highest risk (e.g., nearing lapse, poor performance, big life events), prompting advisors where to focus their time for the most impact. These solutions keep humans in control of key judgments while using AI to handle data gathering, normalization, comparison, drafting, and compliance checks, significantly expanding advisor capacity and consistency across the lifecycle..

2 / Concept Development and User Interaction / Advisor Hub AI

Advisor Hub is a human-centered, AI-enabled "case hub" for life-insurance advisors. It uses document and data-extraction models to pull key fields from compliance documents, carrier PDFs, medical records, and financial statements directly into a structured case record.

Using pre-built connectors and an AI "case assembler," the system, given a client or policy ID, automatically gathers the latest data from CRM, email threads, carrier portals, and cloud folders into a single workspace for the advisor to review, edit, and approve. The goal is to cut proposal and review prep time (for example, from several hours to under 30–45 minutes per complex case) while improving data quality and compliance documentation.

Main interview findings

• Data is fragmented across CRM, email, cloud storage, and carrier portals, forcing extensive manual search, download, and copy-paste work for each case.

• Advisors run multiple carrier illustrations separately, then manually compare options and build consolidated proposals or presentations for clients.

• For annual reviews, staff repeat the same multi-system data gathering and illustration work, then rebuild Excel and PowerPoint reports from scratch.

• Checking recommendations and reports against regulatory rules and internal guidelines is largely manual, increasing the risk of missed disclosures, inconsistent documentation, and failed audits.

• Manual rekeying of values between systems (for example, from carrier portals into Excel or CRM) introduces errors that can affect underwriting decisions and client trust.

• High volumes of small, repetitive tasks (document chasing, formatting, follow-up emails) reduce time available for advice, planning, and client conversations.

• Advisors are generally digitally savvy, mobile, and comfortable with apps, but they want tools that feel like assistants—not black boxes that override their judgment.

•Explicit outcome goals, e.g., "Reduce time spent on data gathering and assembly by 50–75% per case" and "Reduce manual data-entry errors and missing-document findings in compliance audits."

Automation vs. user control

Design principle: Advisor Hub automates data collection, assembly, and drafting, but advisors retain full control over interpretations, recommendations, and client communications.

Data collection and normalization

• Auto-ingest documents and data from integrated systems and extract structured fields into the case record.
• Detect missing or inconsistent fields and generate a checklist of items to obtain or confirm.

Mixed-initiative interaction

• The AI surfaces "next best actions" (e.g., "Run updated in-force illustration," "Request updated financials") that the advisor can trigger or ignore.
• Advisors can ask natural-language questions ("What's changed since last review?") and get draft summaries they can refine.

Case assembly, review, and editing

• Draft side-by-side comparison tables and narrative summaries from multiple carrier illustrations and existing policy data.
• Pre-populate review decks and proposal templates with charts, key metrics, and standard disclosures, leaving advisors to tailor the story.

Ongoing case and compliance updates

• Track key dates (policy anniversaries, term conversions, premium changes) and generate reminders and draft outreach messages.
• Monitor for regulatory or product-rule changes (where data is available) and flag impacted cases for advisor review.

Additional automations to consider

• Intelligent document routing: auto-file incoming client and carrier documents to the correct client/case and tag them (e.g., "APS," "lab result," "policy change form").
• Smart task queues: prioritize advisor worklists by lapse risk, upcoming reviews, and missing-requirements impact.
• Template selection: automatically choose appropriate proposal/review templates based on product type, jurisdiction, and client segment.

User control mechanisms

- Explicit “review and approve” steps before any AI-generated content is sent to a client, carrier, or CRM.
- Per-workflow automation settings (e.g., “auto-draft emails but never auto-send,” “suggest but never add products/riders”).
- Fine-grained edit controls with clear labels on AI-generated sections so advisors know what they need to scrutinize.
- Easy “undo/revert” for AI changes to preserve a sense of safety and control.

Trust plan: avoiding overreliance and algorithm aversion

Transparency

- Source labeling: every field and sentence carries a provenance tag (e.g., “Carrier PDF,” “Salesforce,” “Advisor edit,” “AI summary from sources X/Y”).
- Inspectable logic: when the AI highlights a data point or recommends an action, it shows a short explanation (“Flagged because premium increased by 20% since last review”).
- Clear role framing: in the UI and training, the system is always described as a copilot for licensed professionals, never as an automated underwriter or decision-maker.

Feedback loops and reviews

- Simple in-product feedback (“mark correct/incorrect,” “needs review”) that is logged and used to improve future suggestions.
- Case-level “confidence indicators” that highlight where data is incomplete or low-quality so advisors know where their attention is most needed.
- Regular review reports for compliance and risk teams (e.g., “Top types of AI corrections this quarter”) to tune guardrails.

Avoiding overreliance

- Mandatory checkpoints in high-stakes steps (suitability reasoning, final recommendation language, replacement analysis).
- Prompts and micro-copy that reinforce verification (“Review these numbers against carrier illustration before sending”).
- Training and onboarding that explicitly cover limitations, typical error patterns, and good verification habits.

Avoiding algorithm aversion

- Start with incremental automation: launch first as a “case prep assistant” that only drafts and organizes, then gradually add more active suggestions once trust is earned.
- Show value transparently (e.g., “This case: 37 minutes of manual tasks automated”) so advisors see concrete time savings without feeling displaced.
- Give advisors and compliance teams configuration power: they can decide which workflows can use AI drafting and which must stay manual or template-only.

Low-fi Wireframes

A responsive app design with chat bot could allow for an advisor balancing control and automation and allow for the advisor to collect data and review and generate policy reports. If the LLM is trained on insurance policy and can link to records, Salesforce, policy information, product portfolios, industry trends, and regulations, etc. Here are some screens that capture the initial data extraction idea, report generation but could be flushed out much more with report, compliance, policy, illustration, template and proposal usability taken into consideration:

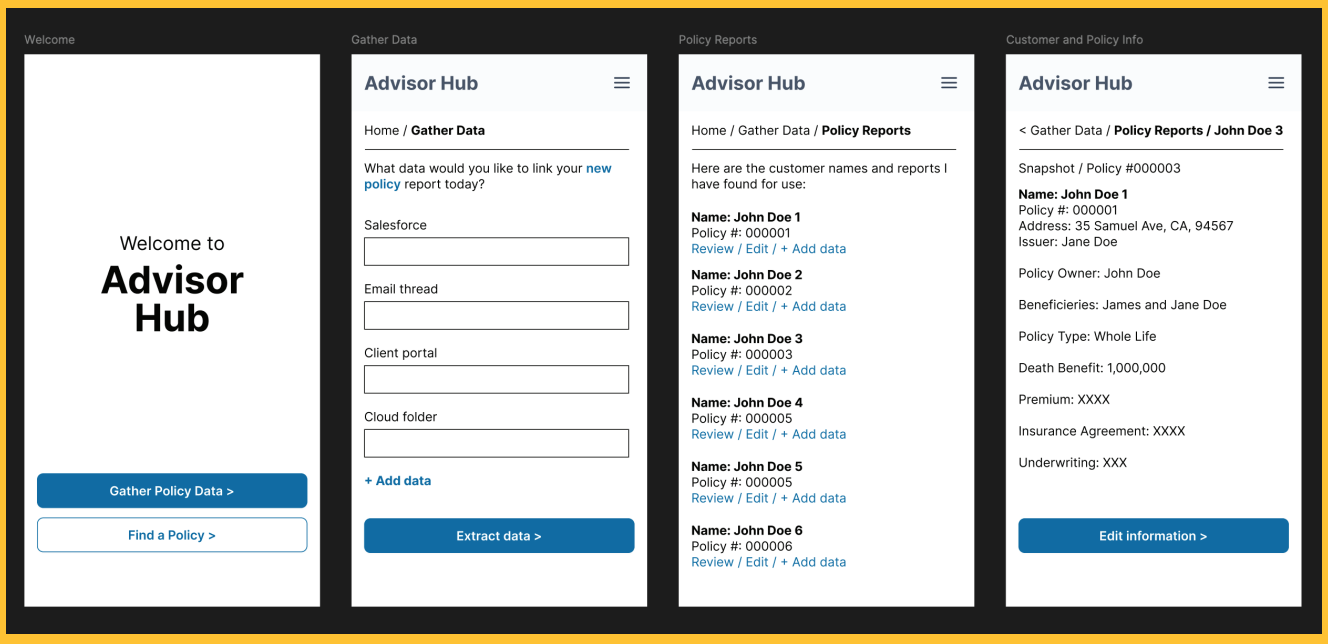

3 / Ethics and Societal Impact Review

There are key ethical principles at stake that need guardrails in the life insurance business considering personal data, advisor bias, legal regulations, data manipulation, current information, underwriting and transparency. Here are some risks that can arise with a product like Advisor Hub AI:

Risks

- Autonomy and dignity: Advisors and clients must keep real decision power, not just decorative sign-off.
- Justice and non-discrimination: Data aggregation and pattern-finding must not quietly encode or amplify bias in coverage or product design.
- Transparency and accountability: It must be possible to explain, challenge, and correct how a recommendation was formed; responsibility cannot be diffused into "the system."
- Privacy and proportionality: Data collection should be minimal and clearly consented, not opportunistic surveillance across every digital trace.

Strategies to mitigate risk

Strategies that mitigate risk can rely on concepts such as drafts, metrics, traceability, questionaires, explanations, fairness governance, human decision and augmentation, user control and training. Here are a few ideas of how to mitigate risk for Advisor Hub AI:

Lock in human primacy for recommendations

- Advisor Hub drafts inputs and comparisons, not final product choices by default.
- UI explicitly distinguishes: "AI-suggested inputs," "Advisor judgment," and "Carrier decision."
- Metrics reward quality (client outcomes, complaint rates), not just speed or "alignment with AI."

Radical traceability and explainability

- Every recommendation must be accompanied by a clear rationale panel: which data sources, what rules or models were applied, and which constraints (client goals, risk tolerance, regulatory rules) mattered most.
- Generate an audit trail that regulators and clients can read: who changed what, when, and why—AI actions and human edits both tagged.

Data and scope limits

- Restrict data sources to those clients reasonably expect (existing financial and insurance records, explicit questionnaires); avoid scraping or opaque third-party behavioral scores.
- Publish a short, human-readable data-use explanation at the point of advisor and client onboarding.

Bias and fairness governance

- Periodically audit Advisor Hub outputs for patterns of proxy discrimination (e.g., by ZIP code, income bracket, protected classes) using standard fairness tests.
- Build tooling for compliance to simulate "counterfactual" clients and ensure similar profiles get similar treatment unless a clear, documented risk factor justifies a difference.

Design for resistance, not only compliance

- Make it easy—and socially legitimate—for advisors to override AI suggestions with structured reasons ("client explicitly prefers lower premium," "other household coverage not in system").
- Have leadership and UX copy emphasize: "Advisor Hub proposes; you decide," and reflect that stance in performance reviews.

Education and shared literacy

- Train advisors and compliance teams on how Advisor Hub works, its limits, and how to spot overreliance or blind spots.
- Provide clients with a short explanation when AI tools contributed to their proposal, plus an invitation to ask questions.