



Quotable: A Search Engine for User Research - Final Business Analysis

Recordings to actionable insights in minutes instead of weeks

Team:

Robert Zhang - robertzh@wharton.upenn.edu | Dapu (Ben) Wang - bendwang@wharton.upenn.edu
Jong Min Choi - jongmin@seas.upenn.edu | Yiheng (Intel) Chen - yihechen@seas.upenn.edu
Cindy Jiang - xinyuji@seas.upenn.edu | **Advisor:** Michael Kearns - mkearns@cis.upenn.edu

Quotable Executive Summary: Conducting user interviews make up the backbone of a UX researcher (UXR)'s job. However, synthesizing insights from user interviews is a time consuming and highly manual process. On average, a single 45 minute interview transcript requires at least 3 hours to properly synthesize. Not only do UXRs find the manual parts of synthesis tedious, manual synthesis reduces the quantity and quality of research output. ***Quotable is a search engine for user research recordings that allows UXRs to get from raw recordings to actionable insights in minutes instead of weeks.***

Our Team's Research: Before we started building, we wanted to make sure we were solving a real pain point. Hence, we spent most of our fall semester conducting rigorous user research. Our candidate user personas to interview were product managers and UX researchers. We operationalized an interviewee recruiting system that enabled us to effortlessly go from cold email to interview. We ended up interviewing **14 product managers** and **12 UX researchers**, amounting to over **17 hours** of recorded user interviews. Not only did this research give us confidence in our product, the long process of synthesizing our own interviews gave us first hand experience in the problem we are solving. Given the volume of our research, **every statement and product feature in this report can be backed up by quotes from users in our target market.** Let's dive in.

The Problem

A UX researcher's (UXR) job is to provide answers to important product questions. Internal stakeholders present UXRs a set of questions they want answered. The UXR then conducts user interviews, and converts interview information into a report with answers backed by evidence. Rinse, wash, repeat.

The problem is, synthesizing insights from user interviews is a time consuming and highly manual process. For every 45 minute interview conducted, UXRs have to spend at least 3 hours finding insights. A standard 2 to 4 week user research cycle contains between 5 and 20 interviews. UXRs have to spend at least half of that time, i.e., 15 to 60 hours per research cycle, conducting interview synthesis.

Where does all that time go? The bottleneck holding up synthesis is the need to manually search for and paraphrase quotes from interview transcripts (*Appendix 1*). **This problem has 3 tangible effects:**

- 1) Manual synthesis increases research turnaround time by adding weeks of synthesis, reducing the relevancy and quantity of research output.
 - 2) Manual synthesis discourages thorough data exploration, reducing the quality of research output
 - 3) Manual synthesis feels tedious for researchers, who would much rather focus their time on high level activities, like facilitating interviews, discovering themes, and mapping user insights to recommendations.
-



The Solution

Value Proposition: Quotable is a search engine for user research recordings that allows UXRs to get from raw recordings to actionable insights in minutes instead of weeks.

Mission: To help UXRs *synthesize as much information as possible, as easily as possible, as accurately as possible* so they can *create compelling, data-backed reports for stakeholders (Appendix 2)*.

Product Workflow: Quotable follows a simple workflow that fits what UXRs are used to (*Appendix 3*)

- 1) **Transcribe:** Users can create a Quotable project, and upload video or audio files from their interviews. These files are automatically transcribed and paraphrased into clear wording for users to skim on-demand.
- 2) **Search:** Most UXRs are accustomed to filling out evidence in a spreadsheet, where columns are questions or themes, and rows are interviewees. Quotable presents information in this same format. However, instead of manually scrubbing through transcripts to fill out their spreadsheet, users can type questions or keywords into our search bar. For each query, Quotable's search engine finds relevant quotes and displays them by interviewee as a Quotable spreadsheet column.
- 3) **Summarize:** Users can click a column or row in their Quotable spreadsheet to summarize. Quotable generates a summary of the quotes under that selected interviewee or query..
- 4) **Explore:** By reading discovered quotes and summaries, users can dig deeper into evidence with follow up searches. As they ask more questions, they unearth even more insights. Finding high level themes (i.e., what questions to ask) is the heart of UXR. *We just make finding evidence for those themes instantaneous.*
- 5) **Export:** Users can export video clips of discovered quotes to add into their reports.

Product Technology:

- 1) **Audio to transcript ML pipeline:** Converts the raw audio of meetings into a transcript with speaker roles identified. **Progress:** We've built an ML pipeline that takes in an audio file and outputs a transcript, with speaker roles and timestamp identified for each word. Using OpenAI's Whisper model as a starting point, we added a multitude of other ML models + heuristics to improve functionality and performance. **The result:** a state-of-the-art transcription model that beats anything on the market right now. We used our model to transcribe our team's interview recordings. *It seriously kicks ass.*
- 2) **Transcript Paraphrasing ML pipeline:** Paraphrases a raw transcript into clear verbiage. **Progress:** Our model is great at this. It's so good in fact, that multiple professors and potential investors have advised us to file a provisional patent (which we are heavily considering).
- 3) **Q&A ML pipeline:** The core of Quotable is its search feature - the ability to accurately and automatically find quotes that answer questions/keywords. We plan to use OpenAI's embedding model to do this. We will also have to fine-tune and experiment with various heuristics to get it working up-to-specifications. **Progress:** we've played around with it and have seen good results. We hope to make continued progress over the summer.
- 4) **Web App:** Outside of these ML models, Quotable is a pure SaaS solution that comes with the usual host of full-stack technology. **Progress:** we've deployed our app on tryquotable.com. You can now sign up, upload audio files for transcription/paraphrasing, and search transcripts. The proof of technology is complete, but we hope to build this into a full product over the summer.



Engineering Innovation:

Quotable is the first time large language models (LLMs) have been applied to this problem domain (user research). In doing so, we've built a host of custom machine learning pipelines that work well (see above) tailored specifically for user researchers.

The Business

Customer Segments: Quotable plans to use UX researchers as its beachhead market, due to their high exposure with interview transcripts. UX researchers are not constrained to those with "UX researcher" in their role title. Since user interviews are a core part of product discovery, product managers, UX designers, and startup founders also conduct user interviews as a central part of their job. In general, our best customers are people where synthesis takes up a large portion of their work week, who dislike the manual searching/paraphrasing process, and would prefer to spend that time on other more engaging tasks.

User Acquisition Strategy: Quotable is B2B2C, selling to UXRs. We plan to use a bottom up sales strategy - pitching individual UXRs on a free trial so they can later convince their managers to free up budget for full capabilities. In this sense, Quotable is designed to have minimal onboarding and learning curve: UXRs can just turn on relevant integrations and get started. Everything else is automatic.

Other Product Stakeholders: To make onboarding easy, no other stakeholder needs to be involved for UXRs to use Quotable. However, two groups of stakeholders will have to be involved in general:

- 1) The buyer: most UXRs are individual contributors with no direct say in their team's budget. They would need to get their manager's buy-in to free up the budget for Quotable. While our value to UXRs is to save them hours of tedious work, our value to the buyer is showing Quotable can empower UXRs to produce better research in higher quantities.
- 2) Client company legal depts: we need to ensure transcripts on Quotable are encrypted and secure.

Market research (TAM, sizing the problem): Our beachhead market of pure UX researchers is relatively small: there are only 40k of them according to LinkedIn Sales Navigator. However, the problem we are solving for this segment is a potent one. Our user interviews indicate these UXRs spend on average 45 hours per month on synthesis. Since most of this time is spent on tedious search/paraphrasing, Quotable can save them at least 35 of those hours. At an average wage of \$60/hr, that's \$2.1k in lost productivity per UXR per month, or **\$1B in lost productivity per year**.

Outside of our beachhead market, there are over 500k+ UX designers, product managers, and startup founders in the United States, and millions around the globe. All of these roles conduct user interviews as a central part of their job (our research of PMs indicate an average of 1 - 2 interviews per week), and would have a need for Quotable's rapid synthesis function. **This amounts to an additional \$5.25B in lost productivity per year in the US alone.** For comparison, Dovetail, a somewhat popular UXR tool (see below), has 100k+ users, and is currently valued at \$960 M.

Competition (Appendix 4): There is no tool out there that allows for automatic synthesis of meeting recordings (and especially no such tool that is purpose-built for UXRs). However, there are tools in



adjacent problem domains. There are **transcription** tools, like Otter.ai, Zoom, Teams, and Rev. For **synthesis**, UXRs use spreadsheet software (like Google Sheets or Airtable) to organize insights, and Miro or Figma for visual affinity mapping. Finally, there are “**end-to-end UXR**” tools, led by Dovetail. They offer transcription, allow UXRs to find and affinity map quotes, and embed quotes as video reels into shareable reports. However, synthesis in these end-to-end tools is fundamentally manual (they lean heavily on users building up tagging systems), have a significant learning curve (discouraging for many UXRs), and require full team buy-in due to an extensive onboarding process.

Revenue Model: We will use a per person tiered pricing model based on compute usage: simple, standard, scale. Pricing is currently fluid, but we believe \$20 - 100/user/month floor price is a reasonable fee, as that is in line with other UXR solutions (like Dovetail), and hence industry expectations.

Costs: Our primary costs come down to the compute needed to run our ML models. The heaviest of these models is transcription, which we only need to run once per uploaded interview. Hence, per-user cost is mostly a function of # interviews uploaded. Additional Q&A and summarization models add variable cost proportional to the number of study questions, but that cost should not be substantial.

Why Now: The power of large language models toward search has captivated the world with the release of ChatGPT and similar internet search products. Not only is this a rapid trend, but the industry (i.e, potential users) in general will have a better understanding of why good search can save time. This will save us CAC in terms of educating the market.

Why Us: The Quotable team is perfectly suited to tackle this problem. Ben, our Head of ML, is an OpenAI researcher who built and trained a copy of GPT-3 from scratch. Jong Min, our Head of Software, is a PennLabs lead highly experienced in full stack development. Intel, Cindy, and myself are all Product Managers who have deep personal experience with this problem.



APPENDIX

Note 1: The Interview Process, Step by Step

Excluding interview planning and recruiting, which are out of our scope, we can generally break down the user research cycle into **3 universal steps**: data collection, data synthesis, and report building.

The **first step** is raw data collection. UXRs always ask to record their interviews, and it is rare for users to reject that request. These recordings are then transcribed by software post-interview. Many UXRs also take live interview notes, either by themselves, or more rarely by asking a colleague who is available to help. Even if notes are taken, interview transcripts are almost always necessary because notes miss details or misquote users. Hence, the output of this step is a set of interview transcripts, and accompanying notes. We believe this step is mostly a solved problem. Many UXRs still struggle with getting high quality transcription, and have to spend significant time fixing transcription errors. However, state of the art transcription companies exist (Otter.ai, Dovetail). This is just a matter of adoption.

The **second step** is data synthesis. After a round of interviews, UXRs are typically left with a set of interview transcripts. UXRs must pull useful data from these transcripts. Pulling data is split into two tedious activities:

- First, searching through transcripts for useful data takes hours. Transcript dialogue is logically meandering and filled with repetition, making transcripts impossible to skim. Researchers have to spend lots of mental energy and time simply deciphering meaning from messy verbiage.
- Second, once the researcher finds a useful quote, they need to spend significant time paraphrasing that quote to be human-understandable and report ready.

Next is affinity mapping. Here, UXRs cluster discovered quotes together to find high-level themes and answers to their given questions. Finding themes and answers from the data is fun and what the heart of UXR is all about. However, when new themes are discovered, they often need to be explored further. This requires more searching of evidence back in the transcripts, in a slow, tedious, iterative process of discovery.

Finally, once the UXR believes they have the answers, and ample evidence to back up their conclusions, they put together a report to share their findings with stakeholders. These reports typically cite paraphrased user quotes, include snippets of interview video, and tie discovered insights to actionable recommendations, to maximize project impact.

Note 2: Product Development North Star (Vision)

Our mission is to help UXRs *synthesize as much information as possible, as easily as possible, as accurately as possible* so they can *create compelling, data-backed reports for stakeholders*.

- *As much information as possible* => make it easy to funnel research of all formats into Quotable. Maximize integrations and native research inputs (native recording/transcription, native feedback portal, etc)
- *As easily as possible* => make Quotable zero-learning-curve for users, minimize onboarding time, increase speed of model inference



- *As accurately as possible* => maximize specificity and sensitivity of search queries, display upfront the most useful information for UXRs with zero/minimal user intervention, preserve objectivity in search results, build trust with users on query results

Quotable's search engine technology can be applied in so many more adjacent fields. In general, our system can help people search any internal source of information. Growth potential is limitless.

Note 3: How Quotable fits into the UXR workflow

Quotable handles the entire synthesis process, from transcription to affinity mapping. UXRs input interview recordings, and export recording snippets, quotes, and summaries to put in their reports.

Quotable does not handle interview facilitation, nor does it host reports. There are so many alternative products that handle both processes, from Zoom, to Teams, to Google Meet, to Notion, to Confluence, to Google Drive, to OneDrive that are often deeply embedded in a company's work processes. Adding those other features will only make Quotable onboarding more difficult.

Note 4: How Quotable differentiates

Unique Attributes: What features set Quotable apart from alternatives?

- Transcript Paraphrasing: transcripts are paraphrased to remove repetition, meandering logic, and other artifacts of speech.
- Transcript Search: instead of manual parsing of transcripts, users can use our search bar to find quotes that answer questions, or are related to themes of interest.
- Research Summarization: instead of parsing long blocky quotes manually to come up with overarching insights, Quotable summarizes quotes to get users started.

Value: What value do the above attributes enable for customers?

All of the 3 features save UXRs the effort and time of tedious tasks:

- It decreases research turnaround time by removing weeks of synthesis, increasing the relevancy and quantity of research output.
- It encourages thorough data exploration, increasing the quality of research output
- It removes tasks that are emotionally draining for researchers, who would much rather focus their time on high level activities, like facilitating interviews, discovering themes, and mapping user insights to recommendations.