MAIN

Chat, Search and Summaries: Smarter Apps with LLMs

Course description

In this practice-oriented project course, students **delve** into the world of large language models (LLMs) to design and develop intelligent applications. Working in teams of two to three people, participants take on the challenge of building applications that leverage LLMs for tasks such as intelligent document searches, dynamic user interactions, and content summarization.

The course spans two months of continuous development and offers a comprehensive introduction to software development. While LLMs form the core of every project, the focus is on integrating and extending their functionality through the following aspects:

- Front-end development: Designing intuitive user interfaces for interacting with the LLM-powered backend.
- Backend development: Architecting robust systems for managing APIs, databases, and application logic.
- API design and integration: Seamlessly connecting LLMs with external systems to ensure efficient data flow and real-time processing.
- Scalability and Deployment: Consideration of performance and hosting aspects to prepare applications for real-world deployment.

Throughout the course, students work in collaborative design and development cycles, receive regular feedback, and iteratively optimize their projects. At the end of the course, teams will present a fully functional application and have a solid understanding of how LLMs can revolutionize modern software solutions.

This course is designed for students who want to expand their skills in developing innovative software while exploring the potential of LLMs.

Next Steps

- ☑ ~~Choose an app idea below or propose one yourself (until Apr 23, 2026)~~
  - ☑ ~~Your name should only show up in ONE project~~

- [x] ~~Make Whatsapp group / Discord / etc- share your project plan below (until Apr 30, 2026)~~
- [x] ~~Submit your 5-min mid-term video (by Jun 3) — we watch & discuss on Jun 4. Get feedback and adapt your project plan (Jun 4, 2026)~~
- [ ] Final presentations — view videos and discuss in class (Jul 16, 2026)

People / Project Directory

Names (per group)	Project (+ link to project page)	Chat handles	Mid-Term Presentation Video	Final Presentation Video
Edvin Perfundi Ferhan Zeynep Aydın Mustafa Can Vera Tereshchuk	Style assistant Project doc		StyleSync.mp4	StyleSync Final
Eleonora Baran Vu	Financial Management & Investment Advisor FinSight_Project_Doc_v2.d…		Finsignt_midterm_presenta…	FinSight Final
Syed Faraz Chiara	Hypothesis Testing Assistant Hypothesis Testing Assista…		Statistical Hypothesis Testi…	Hypothesis Testing Assistant F…
Tanya	AI Menu: MenuMind Project doc		Presentation	Final Video MenuMind Try it as Web app! https://menu-mind-tawny.vercel.app/ Try an Android app 🙂 https://play.google.com/apps/testing/com.menumind.mobile (contact me for details)
Daniel Yu Leonard Goldmann	Game: I am a Robot Project Doc	Amanda Dori		

Nadir Bašić Luka Radulovikj Eric Germersdorf	Flat Chat - Berlin Real Estate (AI) Agent Flat Chat - Project Doc - S…		Video	Final Video Flat-Chat Try it out! https://flatchat.cc/ (ask for password first)
Khurram Hamna Urooj Irfat Niso	ADHD Assistant Project Doc		midterm: Video	Link 2 Telegram bot: @secondbrain01_bot Have a chat, it's live. Github: https://github.com/uroojkhan01/SecondBrain_CognitiveSystem
Siji Mehri	Official Letter Assistant Link: Project Doc - Completed.d… GitHub: https://github.com/Mehmar-11/Official-Letter-Assistant		Video Presentation	Letter_Assistant_Final_Prese… Try it out! https://official-letter-assistant.vercel.app/
Adnaan Ahmed Avadhoot Sandeep Raikar Shameem Khan Botan	Abstract Syntax Tree Retrieval Assistant Project Doc		Presentation	Presentation Results
Benedikt Peterson Alireza Amini	Smart Form filler (browser extension with LLM API) Link: Project-Doc		Smart-Form-Filler-Demo.mp4	Smart-Form-Filler-Final-Version
Anna Benedict Eren Kaleem	Hired: Career Planning Agent Link: Project-Doc		Presentation Link	Hired_Final_Presentation.mp4 Click here to download the latest version for free!

	Wearables and Well-being			
Athulya wilson Shyby Joseph Sandra jojo	LLM-Based Parametric 3D Generator link:project doc		Presentation Link	3D model generator final Click here :https://frontend-mocha-three-73.vercel.app
Saad Waseem Jingren Dai	Intent-Based Passive Search Agent App Project Doc		Presentation Link	Final Presentation
Saloni Kothari	Intelligent daily task planner			
Nilüfer Cagla Yigit	CineMapper: Cross-domain Movie Recommender		Video presentation link	Final Presentation
Leon Koch David Glaser	cs-modulio.com CS Modulio Doc		Video Presentation	Final Presentation

Projects / App ideas

Games Against LLMs

Develop interactive games where the user competes against or cooperates with an LLM, testing the model's strategic reasoning and creativity. (e.g. manipulate AIs in “social” interactions)

Educational Tutors

Build an interactive learning platform that uses LLMs to provide personalized tutoring, generate custom quizzes, and explain complex concepts in an easy-to-understand way (e.g., learn a topic by talking to an LLM).

Language Learning App (Rare Languages)

Design a conversational language learning app that acts as a native speaking partner for rare languages (like Danish). It corrects grammar, expands vocabulary, and simulates real-life scenarios.

Smart Travel Assistant

Develop an application that helps users plan itineraries, discover local attractions, and manage bookings by conversing with an LLM trained to act as a seasoned travel guide.

Financial Management & Investment Advisor

Create a tool that digests financial news, crypto market trends, and investment data to provide summarized insights and educational financial advice to users.

Hypothesis Testing Assistant

Build a tool for researchers that uses LLMs to brainstorm, formulate, and critically evaluate scientific or business hypotheses against available literature and data concepts.

Medication Interaction & Suggestion Guide

Develop an informational assistant that helps users understand their prescribed medications, checks for potential interactions, and explains side effects (with appropriate medical disclaimers).

Microfacial Expression Recognition

Combine vision models and LLMs to analyze brief facial expressions from video input, providing users with insights into emotional reactions during interviews, negotiations, or presentations.

Sales Lead Generation Bot

Create a tool that scours the web or social media to identify potential B2B leads, automatically drafting personalized outreach emails or LinkedIn messages tailored to each prospect's background.

LLM-Powered Operating System Interface

Design an experimental, conversational interface for an OS where users can manage files, run scripts, and configure system settings entirely through natural language commands.

Business Growth Management

Develop an AI consultant that analyzes business metrics, market trends, and customer feedback to suggest actionable growth strategies and optimize operational efficiency for startups.

Backtesting Engine for Business Plans *

Build a simulation engine that subjects user-submitted business plans to hypothetical market scenarios, competitive responses, and economic shocks using LLM-generated roleplay and analysis.

Intelligent Daily Task Planner

Create a smart calendar application that not only organizes to-do lists but dynamically schedules tasks based on priority, user energy levels, and estimated completion times. Integrate data from smartwatches and wearables with an LLM to provide holistic health insights, daily motivation, and personalized lifestyle adjustments for improved physical and mental well-being.

Work-Life Balance Policing

Develop an application that monitors work habits, alerts users when they are overworking, and enforces breaks through playful or strict LLM-driven interventions.

Automatic Context-Aware Subtitles

Build an application that uses speech-to-text combined with LLMs to generate highly accurate, context-aware, and translated subtitles for videos or live streams in real-time.

Social Media Personal Branding Assistant

Create an AI manager that helps professionals curate their online presence by suggesting post ideas, drafting engaging content, and analyzing audience engagement to build a strong personal brand.

BigTalk Topic Determinator

Design a conversation-starter app that skips the small talk and generates deep, philosophical, or controversial "big talk" topics to spark meaningful and intense discussions among friends.

Doctor's Appointment Preparer

Develop an assistant that helps patients prepare for medical visits by organizing their symptoms, generating a list of questions to ask the doctor, and summarizing their medical history.

Smart Form Filler

Build an intelligent tool that stores a user's personal or business profile and automatically parses and fills out complex web or PDF forms, saving time on repetitive applications.

Skin Cancer Detection Web-App

Build a full-stack medical web application that allows patients to photograph skin lesions and receive a preliminary diagnosis by using a computer vision model, with LLM-generated explanations, while giving doctors a dashboard to manage patients, review AI predictions.

Berlin AI Assistant

Scrape the web (+ more?) for apartment listings. Build an intelligent web app for listed apartment interactions. Show on a map, show proximity to public transport, amenities… Chat with an assistant and help choose and apply to apartments…
Other ideas:

- Admin assistant (e.g, visa information / registration, apartment Anmeldung, …)
- Events in Berlin
- Help find UNI Courses in Berlin + Potsdam

Career planning Agent

A career planning agent that automates the entire career planning and application process end-to-end. It researches open positions, generates personalized cover letters and CVs, sends applications via Gmail, and fills out web forms on company portals. It also provides you the info about the company, gives an insight into possible salary range.

Game: I am a Robot

Name not final.
The idea is to have a game where one or multiple human players try to pass as an AI. Maybe we'll do it the other way around or just do both as different gamemodes.

Style Assistant

Web-based AI styling suggests your outfit for certain occasions and situations. You can upload your clothes once and the app handles the rest, like suggesting daily outfits based on your mood, the weather, and your schedule.

Official Letter Assistant

Official Letter Assistant helps people understand German official letters by turning complex bureaucratic language into clear, actionable guidance: what the letter says, what matters most, what deadline exists, what the user must do, and what uncertainty remains.

Example Project Doc

Example Project Doc

0 . Cover

Project name: Smart Literature-Review Assistant
Team members: Alice Müller (PM), Bob Singh (FE), Carla Rossi (BE)
Date / version: 2026-05-06 – Draft v0.2

1 . Problem Statement & Vision

One-sentence pitch – “Help researchers query hundreds of papers like ChatGPT.”
Why now? / Stakeholders – Pain points, target users, success criteria.

2 . Scope

Must-have (MVP): Upload PDFs, ask Qs, get answers
Nice-to-have: Export summary to Word
Stretch goals: Citation graph

Explicitly list what is out of scope to prevent scope creep.

3 . User Stories

“As a PhD student, I want to drop 20 PDFs and ask “Compare methods”, so that I can skip manual reading.”

… add more

4 . System Overview

- Architecture diagram (boxes & arrows).
- Key components – FE, API, vector DB, LLM provider.
- Data flow – upload → chunk → embed → retrieve → generate.

5 . (Optional) Tech Stack & External Services

…

6 . Project Plan & Timeline

Week	Milestone	Owner	Deliverable
1	Repo, CI/CD, skeleton FE	Bob	URL + README

2–3	PDF-ingest pipeline	Carla	/api/upload working
4	First Q&A demo	Team	Live demo
5	UX polish, tests	Alice	Usability test results
6	Stretch goal pick	—	PR merge
7	Load test & deployment	Team	Prod URL
8	Final presentation	Team	Slide deck

7 . Roles & Responsibilities

Role	**Person**	**Key Duties**
Product / PM	Alice	Scope, timeline, stakeholder comms
Front-end Lead	Bob	UI, accessibility, tests
Back-end Lead	Carla	APIs, DB, DevOps

Define a lightweight decision-making process (e.g., “two approvals per PR”).

8 . Risk Register

Risk	**Probability**	**Impact**	**Mitigation**
OpenAI quota limits	Medium	High	Cache embeddings, queue requests
PDF parsing errors	High	Medium	Fallback to OCR, log fails
Exams overlap	Medium	High	Plan buffer week 5

9 . Evaluation & Demo Plan

- Success metrics – answer accuracy @ Top-1, latency ≤10 s, user SUS ≥70.
- Demo script – step-by-step flow to show during final presentation.
- Backup video in case Wi-Fi fails.

10 . (Optional) Appendices

- Reference material, prompt library, style guide links.
- Meeting minutes index.
- Glossary of acronyms.

Evaluation Criteria

Evaluation Criteria

For grading your projects I will calculate a weighted sum of points you earn along a handful of evaluation dimensions (weight in brackets):

1. **Core functionality and correctness (10)**- does the end-to-end chat/search/summary workflow work reliably?
2. **LLM integration quality (7)**- how thoughtfully are retrieval, prompting, grounding, cost, and safety handled?
3. **Engineering practice (5)** - code clarity, tests, CI/CD, and deployability.
4. **User experience (5)** - clarity, accessibility, and responsiveness of the interface.
5. **Documentation and presentation (10)** - completeness of the project plan, API docs, and demo narrative (mid-term video + final presentation).
6. **Innovation and ethical considerations (5)** - originality, real-user value, and care for privacy or bias.

Each dimension carries a defined point weight, with bonus credit for finished stretch goals and polished extras, giving students a transparent target while still rewarding creativity and thoroughness.

Previous projects - Summer 2025

Previous projects - Summer 2025 (check out their project pages!)

Smart Literature-Review Assistant

What it does:
Researchers drop a pile of PDFs and ask natural-language questions; the app produces threaded answers plus a “living” summary.

Core LLM Tasks:
QA over papers, summarisation, extraction of key ideas, verification of

Tech / Domain extras & datasets:
PDF-to-text pipeline, LLM Backend, Frontend

Stretch goals:
Summarisation

People:

Smart Forms Filler

What it does:
Stores a user’s “master profile” (e.g., personal / company details, CV snippets, standard disclosures) and automatically maps those data points to any web or PDF form. Ideal for startups repeatedly filling investor questionnaires, job-seekers applying to many portals, or researchers submitting grant applications.

Core LLM Tasks:
Field-label understanding and synonym mapping (“VAT ID” ↔ “USt.-ID.”), entity extraction from profile text, gap detection (“You haven’t provided a D-U-N-S number—please supply or skip”), and generation of context-tailored answers where free-text is needed (“Describe your company in ≤ 500 chars”).

Tech / Domain extras & datasets:
HTML form scraper (BeautifulSoup) and PDF AcroForm parser, schema-less profile store (JSON + embeddings in a vector DB for fuzzy lookup), browser extension or Electron wrapper for one-click auto-fill, open-source form corpora (e.g., SEC EDGAR exhibits) to pre-train field synonyms.

Stretch goals:
Secure local encryption vault; learning loop that refines mappings from user corrections; multilingual form support; compliance checker (e.g., flag GDPR-related fields); e-signature integration for final submission.

Doc: Smart Forms Filler Project

People:

Context-Aware Customer-Support Dashboard

What it does:
For small e-commerce shops: agents see auto-drafted replies that cite company docs and past tickets.

Core LLM Tasks:
Ticket summarisation, retrieval-augmented answer drafting, tone adaptation.

Tech / Domain extras & datasets:
Ingest FAQs, manuals, Shopify API; role-based UI (agent vs. manager).

Stretch goals:
Real-time sentiment monitor; self-service customer chat widget.

People:

Campus Navigator Chatbot

Wh it does:
Students ask “Where’s a free study spot with power sockets?” Bot blends static campus data with live occupancy feeds.

ZCore LLM Tasks:
Multi-turn chat, geospatial reasoning, condensed answer cards.

Tech / Domain extras & datasets:
Public campus map, room-booking API, IoT occupancy feed (or mocked data).

Stretch goals:
Voice interface; AR mini-map via WebXR.

Meeting Minutes & Action-Item Tracker

What it does:
Upload call recordings; the app transcribes, tags speakers, and outputs bullet minutes plus per-person tasks.

Core LLM Tasks:
Long-form summarisation, speaker attribution, action-item extraction.

Tech / Domain extras & datasets:
Whisper / Deepgram for ASR, calendar API, Trello / Jira webhooks.

Stretch goals:
Topic-timeline view; confidence heat-map of transcript.

People:

Doc: Meeting Minutes & Action-item Tracker

Explain-My-Codebot

What it does:
Devs query huge codebases: “Where is the OAuth flow?” or “Summarise this function.”

Core LLM Tasks:
Retrieval over code embeddings, explanation generation, diff summarisation.

Tech / Domain extras & datasets:
Git repo ingestion, tree-sitter parsing, code embeddings (OpenAI code models).

Stretch goals:
PR reviewer that suggests tests; natural-language-to-code snippet search.

People:

Personalised News Digest & Q&A

What it does:
Users choose interests; receive daily digest emails and can chat follow-up questions (“Why does this matter for EU policy?”).

Core LLM Tasks:
Multi-source summarisation, relevance ranking, context Q&A.

Tech / Domain extras & datasets:
RSS crawler, NewsAPI, user-profile embeddings.

Stretch goals:
Bias-detection slider; concept-relationship graph.

People:

Multilingual Travel Planner

What it does:
User chats in any language to plan a trip; the app merges guides and live data, returning itinerary and budget sheet.

Core LLM Tasks:
Multilingual chat, tabular summary output, cost aggregation.

Tech / Domain extras & datasets:
Amadeus / Skyscanner, open travel guides, currency API.

Stretch goals:
Group-trip mode with shared edits; offline PWA.

People:

Reflective Journal Analyzer

What it does:
Prompts users for daily entries; provides mood summaries, recurring-theme insights, and gentle suggestions.

Core LLM Tasks:
Sentiment analysis, theme clustering, empathetic response generation.

Tech / Domain extras & datasets:
Local-first storage for privacy, optional end-to-end encrypted sync."

Stretch goals:
Long-term trend visualisations; CBT-style prompt library.

People:

Doc: Reflective Journal Analyzer

Ingredient-Aware Recipe Generator

What it does:
Scan fridge items → get recipe chat suggestions that cite sources and explain substitutions.

Core LLM Tasks:
(Optional) vision-to-text, constrained generation, step-wise explanation.

Tech / Domain extras & datasets:
Supermarket price API, USDA nutrient DB, image-upload pipeline.

Stretch goals:
Auto-generated shopping list; AR overlay for cooking steps.

People:

Patent Finder

What it does:
Lets engineers, researchers, and startup founders search worldwide patent corpora with natural-language questions (“Show me battery-tech patents that solve dendrite growth”) or by uploading a draft invention disclosure for prior-art checks. The app returns ranked matches, highlights overlapping claims, and produces plain-English summaries of each patent’s novelty and scope.

Core LLM Tasks:
Semantic retrieval over patent abstracts + claims, long-document summarisation into key claims & inventive step, CPC/IPC code classification, similarity scoring between a user-supplied idea and existing patents, and jargon-to-layman rewriting.

Tech / Domain extras & datasets:
Bulk data feeds from USPTO PatentsView, EPO Open Patent Services, and Googl6e Patents

(BigQuery); ETL pipeline to chunk & deduplicate filings; vector DB (Qdrant | FAISS) for semantic search; claim-structure parser (using spaCy patterns or tree-sitter for XML); front-end with in-document hit highlighting and PDF viewer.

Stretch goals:
Novelty “heat-map” that visually compares a draft claim set to prior art, email alerts for new filings in a watchlist domain, claim-diff view with auto-generated landscape charts, diagram OCR to match technical drawings, and an export to patent-office IDS (Information Disclosure Statement) format.

People:

OpportuneX - Intelligent Job Recommendation Platform for Strategy Consulting

What it does:
A specialized job aggregation platform for strategy consulting roles in Germany. It intelligently scrapes career pages of target consulting firms using an LLM-assisted navigation and extraction system. Extracted job postings are then structured, enriched, and categorized by an LLM for a comprehensive and accurate job listing service.

Core LLM Tasks:

- Assisting web scraper navigation on complex/dynamic career pages.
- Extracting and structuring job details (title, description, location, requirements) from HTML.
- Categorizing jobs by role type, seniority, consulting tier (T1/T2/T3), work location type, required experience, and degree.

Tech / Domain Extras & Datasets:

- Custom web scraping engine (Python, Selenium) with LLM-guided extraction (Google Gemini).
- Data processing pipeline (Python) for structuring, categorization, and deduplication.
- PostgreSQL database for job and company data storage.
- React.js/Typescript frontend for job display and interactive filtering.
- Focus on the German strategy consulting job market.

Goals:

- Launch a specialized platform for strategy consulting jobs in Germany, initially focusing on ~20 top firms, then expanding.

- Achieve superior data quality, comprehensiveness, and freshness compared to generic job aggregators within this niche.
- Deliver a user-friendly interface with niche-specific, intelligent filters.

Stretch Goals:

- **Web Scraper Module:** LLM-assisted navigation, targeted data extraction, scheduling, and error handling.
- **Data Processing Module:** LLM-driven data structuring, categorization (tier, seniority, etc.), and validation.
- **Frontend Module:** Clean job listing display, responsive design, and robust filtering (tier, location, seniority, work type, experience, degree) and search.

People:

Project Plan Word: OpportuneX Project Plan Word (We worked in Word Desktop, in the Word online version or in Google Docs there are sometimes formatting errors)

Project Plan Google Docs: OpportuneX Project Plan Google Docs

For a quick understanding: ChatBot for Project Plan

The problem: Job portals often do not list all vacancies

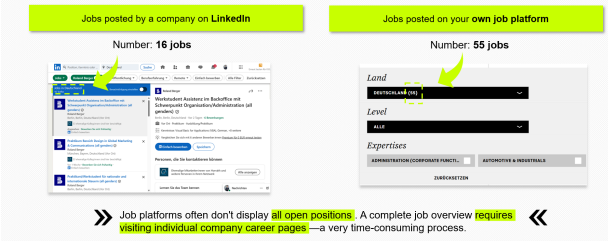

Possible structure of a software project consisting of web scraper , middleware and frontend

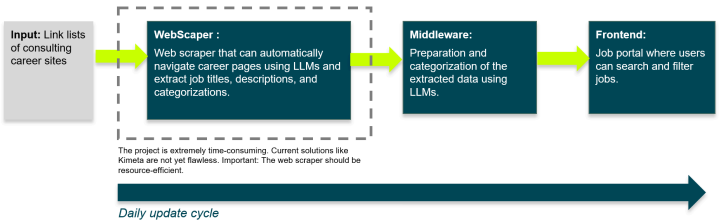

Meal Plan Generation and Tracking, Incentivizing the User - ChallengeX

What it does:
The application would focus on the meal plan generation customized as per the preferences of the choices and the habitability, The meal plan generated will be as per the consultant nutritionist by the LLM.

The major concern which raises up is the motivation for the user to follow down the meal plan generated, so the Multimodal LLM would be accompanied to track down the progress and whether it goes as per the recommended meal plans, which would score the user, and they compete with their friends.

Core LLM Tasks:

1. Meal Plan Generation Planning
2. To Analyze the Meal Plan and Scoring System
3. AI Agents to help the user to get recommended the healthy food of their choice, with the location and the place to get it from.

Tech / Domain extras & datasets:
The Langchain framework would be best to accompanied with the Firebase database to store the LLM generated result and the rendering. For the multimodal LLM the consideration is for the Gemini Flash series as it analyzes the images very well and also has been trained on a lot of diverse data.

Stretch goals:

The product can be extended towards the money lending with the token transaction to generate the meal plan and to follow on which will help to incentivize the user from the pool, which would encourage and motivate to follow on the challenge of the meal plan on the portal.

Proposed by:

Doc: Healthy Meal Plan Platform

Portal: https://challengex-xyz.web.app/

Repository: https://github.com/Hrishikesh332/ChallengeX-xyz

LegalOps Case Intelligence Platform

What it does:
An AI-powered platform built for law firms and in-house legal teams to manage case-specific meetings and documents. Lawyers upload call recordings and legal files per case; the platform transcribes calls with speaker tags, summarizes discussions, extracts action items, and analyzes contracts or filings for key clauses, compliance issues (e.g. GDPR), and legal risks. Everything is stored and searchable by case.

Core LLM Tasks:

- Case-specific meeting summarization
- Speaker attribution and per-lawyer task extraction
- Semantic clause identification in documents and transcripts
- Risk-scoring and compliance summarization (e.g., GDPR)

Tech / Domain extras & datasets:

- Whisper / Deepgram for high-accuracy legal transcription
- SEC EDGAR contracts & other public datasets
- Clause taxonomy tailored to GDPR, IP, indemnity, and more
- Calendar & case management API integrations (e.g., Clio, MyCase)
- Webhooks for Trello / Jira for legal task automation

Stretch goals:

- Timeline view showing legal topics discussed over time
- Transcript confidence heat-map for quality assurance
- Real-time red-line clause drafting based on firm playbooks
- Jurisdiction-specific clause comparison tools
- Full-text semantic search across calls, notes, and files per case

People:

https://docs.google.com/document/d/1SzlcLdLmwDCr6rHbUT3IrBpoPOfcwkOkuJsR5Gnrdb0/edit?tab=t.0

Stylist Assistant

Limit on budget, sizes of clothes, styles they like → assistant finds outfits

People:

Mental and emotional well-being companion

What it does:

Guided chats about mental and emotional well-being.

Analyzes responses to detect if therapy might be recommended.

Generates a report explaining why, highlighting key statements

Suggests coping tips with credible sources.

Not a substitute to therapy or professional service

Core LLM tasks:

Detecting overall sentiment and emotions

Keeping track of key indicator words and quote them for transparency

Suggesting evidence-based tips

Tech / Domain extras & datasets:

Multiple LLMs (GPT4, Llama.. etc.)

Simple tech stack for a web app

Stretch goals:

Implementing a scientific method to base the chats personalized to individual need

People:

Document plan: Mental and Emotional Well-being Companion App

Text-Guided Audio Retrieval and Source Separation using Large Language Models

The goal is to develop a system where a user can upload an audio file and describe in natural language which component(s) they want to hear, for example, *"the vocals only"* or *"remove drums but keep bass and synth."* Based on this input, the system will:

- **Retrieve** a matching pre-isolated audio source from a database if available, or
- **Fallback to separation**, extracting the requested source(s) from the full audio using pre-trained models if no suitable match is found.

The **core focus** of the project is on the **LLM-driven interpretation of user queries** and how they guide the retrieval or separation process. The system will determine which approach to use based on the semantic understanding and confidence of the match.

To keep the emphasis on natural language understanding and decision logic, the project will utilize existing pre-trained source separation models (e.g., Spleeter, Demucs) rather than training new ones.

People:

Text-Guided Audio Retrieval and Source Separation using LLMs

LLM-Based Study Coach:

A platform made for students, by students, aimed at helping to structure self-study, revise concepts efficiently, and test themselves to prepare for exams.

What it does:
A program where students can dump all their study contents (Notes, Slides, Assignments, PDFs, URLs) and the LLM can automatically:

- Generate interpretable study notes and explanations (maybe the user can tell how long they have left to the exam)
- Create flashcards to help retain concepts in memory and revise quickly
- Generate practice questions and quizzes
- Lets students track their progress
- Explains the answer to any questions or clear up any confusions that the student can have about any concept in the slides
- Create a crash course study plan for rapid exam prep

Core LLM Tasks:

- **Summarization:** Summarize uploaded content by topic/section.
- **Key Concept Extraction:** Identify important terms, concepts, formulas.
- **Flashcard Generation:** Create Q&A pairs from key points.
- **Quiz Generation:** Generate multiple-choice or short-answer questions.

Potential Tech Stack:

- **Frontend**: React + TailwindCSS (or Streamlit for faster prototyping).
- **Backend**: FastAPI/Flask + simple PostgreSQL **OR** just use Supabase
- **LLM**: GPT-4 Turbo/Gemini 1.5 Pro (API + Prompt Engineering)
- **Storage**: Firebase or Supabase (for storing user uploads and flashcards).
- **PDF/Text parsing**: pdfplumber, BeautifulSoup for URLs.

Stretch Goals:

- **Adaptive Feedback:** Adjust question difficulty based on user performance.
- **Voice mode**: Reading questions aloud.
- Personalizing the study plan according to how long till the exam.
- Gamified quizzes (with badges or levels, think Duolingo but less murderous owl)

Doc: LLM-Based Study Coach Plan

Group Members: