

sugarlabs



Sugar Labs

Google Summer of Code 25'

Project Information

Interactive AI-powered Chatbot and Debugger for Music Blocks

Length: 350 Hours

Mentor: Walter Bender

Assisting Mentor: Devin Ulibarri

Student Information

Full Name: Amrit Rai

Email: iamamrit27@gmail.com

Github: github.com/retrogtx

LinkedIn: linkedin.com/in/amritwt

X (Formerly Twitter): x.com/amritwt

Preferred Language: Proficient in English

Location: Navi Mumbai, Maharashtra, India

Timezone: IST (UTC +5:30)

Phone: +91 7900064577

Institution: Terna Engineering College

Program: Electronics and Telecommunication

Expected year of graduation: 2026

About

Being into code for a brief period of time, open source is something that affected my life without me even knowing about it. It felt amazing to me that I could contribute to stuff that could be used by people.

My programming journey started back in 2020 when I made ML models with PyTorch, and learnt about core classical ML. It was an amazing experience.

Over time I wished to make the product side of things, so I pivoted to fullstack engineering. This was key as I could ship something daily. With project based learning, I could just build things I wish that existed.

In November 2024, I got a grant of \$3000 by the Solana Foundation for solving inheritance on Solana, a way to prevent dead wallets and pass on crypto to your loved ones. This was a great time for me to get my hands dirty with Rust.

In January 2024, I got into Cal - an open source scheduling software. I've been an intern here for the past one month. The things I have learnt here have shaped me better as an engineer. I write tests, implement features, and review code by other folks from open source. And read code all the time.

On the side, I try to be better at algorithms so the best way to do so is getting hands on competitive programming. Still an 893 rated newbie though.

I try to help others in the community through my small twitter presence.

Tech Stack

I am generally tech stack agnostic, picking up whatever I need to achieve what the problem needs of me. However I am experienced with all these technologies:

Programming Languages: Python, Javascript, Typescript, C++

Core Libraries: NextJs, React, Express, Pytorch, Flask, FastAPI, Drizzle

State Management: Recoil Js, Zustand

Databases: PostgreSQL

Tools: Git, Github, AWS, Docker, CI/CD, Supabase

My tech stack has evolved over the years, when I first started with ML it was mostly about Python, pytorch and other libraries for data science like pandas and matplotlib

For my grant project I worked deeply with Rust, Anchor and Solana/Web3.js

For my day to day projects, it's all about Typescript (or Javascript)

I pick things up as required. But by experience wise, I have listed everything above!

Why Sugar Labs?

When I first started to learn how websites on the internet are built, I started to learn HTML. Then overtime I wished to learn React. My teacher told me to go deeper into DOM manipulation. This is when I got deeper into Music Blocks. This was a great way for me to learn about DOM manipulation, and to see what a full fledged project with vanilla Javascript looks like.

As a fellow music enthusiast, it was nice to see that there is an app out there that helps kids learn better about music.

By diving deeper, I found the way it works was remarkable. Just a huge codebase with javascript running. I started playing around from day one. Figured out what part of this huge might be nonoptimal and started contributing code to solve the same.

Contributions were one part, but I have also been active in the community, helping with small stuff here and there.

I don't just contribute; I often actively participate in the community, sharing ideas and working together on solutions at times.

Issues Raised

[feat]: Add Dark Mode (#4195)

Musicblocks was light mode only.
Dark mode was added in subsequent PRs by the community.



Pull Requests

calling every possible direction from a single const row (merged)

Streamline Drum Name for early return of drum post http match, use object lookup (merged)

move all saved state into a single object (merged)

refactor: simplify note processing logic, remove an empty file (merged)

todo: apply array destructuring (merged)

split notation mappings into separate Octaves (merged)

fix #2630: Add jsdoc style documentation (merged)

Fixes: #4056 correcting highlight problem near nav-bar (merged)

Increasing the size of chord pie menu (merged)

Fixes: #4056 correcting highlight problem near nav-bar (merged)

refactor: use better mapping logic in _setupBlocksContainerEvents (merged)

fix: base64encode error warning, extraction of common logic into a single function (merged)

fix renderLanguageSelectIcon logic (merged)

fix all eslint errors along with base64encode error (merged)

FIXME: Implement pickup measures and multi-voice support in abc.js (merged)

Streamline Drum Name for early return of drum post http match, use object lookup (merged)

use object lookup for convertDuration instead of switch (merged)

fix all lint + Base64Encode errors (merged)

use object lookup for convertDuration instead of switch (merged)

move all saved state into a single object (merged)

used regex to minimize code (merged)

add export statements for BACKWARDCOMPATIBILITYDICT and initBasicProtoBlocks (merged)

Update drum block setup by combining everything into a list (merged)

refactor: Simplify note processing logic, remove an empty file (merged)

<https://github.com/sugarlabs/musicblocks-v4/pull/415> (not merged)

A full list can be found [here](#).

And many instances where I have helped with reviewing and more just by lingering around :D

Previous Contributions in Open Source


A contributor at [Cal.com](#)

When I first started to code back in 2020, I was a machine learning engineer, making stuff here and there. Over time I wanted to work and be at the product side of things. Build stuff that people could use. So I started with fullstack engineering. I also participated at hackathons, and there I came across Cal - a meeting scheduling company that helps in connecting professionals, companies and people in general by managing their time for them. That is the first time I used it.

Then I came to know that the app I use right now is open source. I made some commits, hopped on a call. In the meantime I got a grant project I started working on. Then moving forward in January, I was offered an internship that would start in February. Since then I have made high impact PRs, but also small fixes that would help the team move forward quicker.

It started from February, and will go on till the end of May. I have since then reviewed other PRs from people too. To check every person making a PR is something that makes me better as a programmer because I get to read a lot of code, and how it impacts the rest of the codebase. Got an internship here!

As of 3rd March, one month of internship has passed. These are the number of contributions I have made:



Through this, I get to play around with NextJS, TRPC, Docker, Turborepo and many other technologies that is required at a production grade application

I hope that I can transfer my skill set to build great stuff at Sugar Labs as much as possible.

[Eternal Key](#)

I built [eternal key](#), after pondering over the question that Satoshi Nakamoto (inventor of Bitcoin) wanted the world to have a decentralised financial system. But there are only 22M of them, out of which there's many that are lying around in wallets of people who might not be alive, making the wallets inactive. The BTC is then wasted, right?

Ideally, we should have a way to inherit crypto but we do not. This is odd as after a point, there will be more wasted currencies on the blockchain than the ones in supply over a longer time horizon.

Building on this idea, I began to reflect on how, from the very creators of cryptocurrencies to individuals like us who hold them, there is no clear way to pass these assets down to future generations.

Consider this scenario: What if the original creators of a cryptocurrency, such as Satoshi Nakamoto, were to suddenly pass away? The crypto holdings they accumulated would essentially remain locked in a digital vault, inaccessible to anyone else. But what if there was a desire to pass on these assets? Simply sharing the seed phrase with someone is risky and impractical. Storing the seed phrase in a bank? That could easily lead to theft or loss. So, why not leverage the blockchain itself to solve this problem?

By utilizing the blockchain, we could create a system where ownership and access rights are secure, transparent, and potentially inheritable. This would enable us to not only protect the value of these assets but also ensure that they can be passed on safely, even if something were to happen to the original owner. The question becomes: how can we integrate these principles into the existing blockchain infrastructure to create a secure, decentralized inheritance system for cryptocurrency assets?

My project got funded by the Solana Foundation because it solves exactly this. A way to inherit crypto, in less than 6 clicks you can choose who will inherit your crypto after some period of time passes.

I also wrote a detailed [blog](#) about this!

This project will be important in the coming years.

Some Personal Projects [Related to ML models]

Attention Is all you need

The paper that started it all, an implementation. LLMs that we use day to day, have originated because of this. The core idea of the transformer model is the attention mechanism, which allows the model to weigh the importance of different words in a sentence relative to each other, regardless of their distance. The self-attention mechanism computes a representation of the input sequence by attending to all positions simultaneously. This is in contrast to RNNs or Long Short-Term Memory (LSTM) networks, which process inputs sequentially and struggle with long-range dependencies. By utilizing this attention mechanism, transformers can capture complex relationships within the data more efficiently.

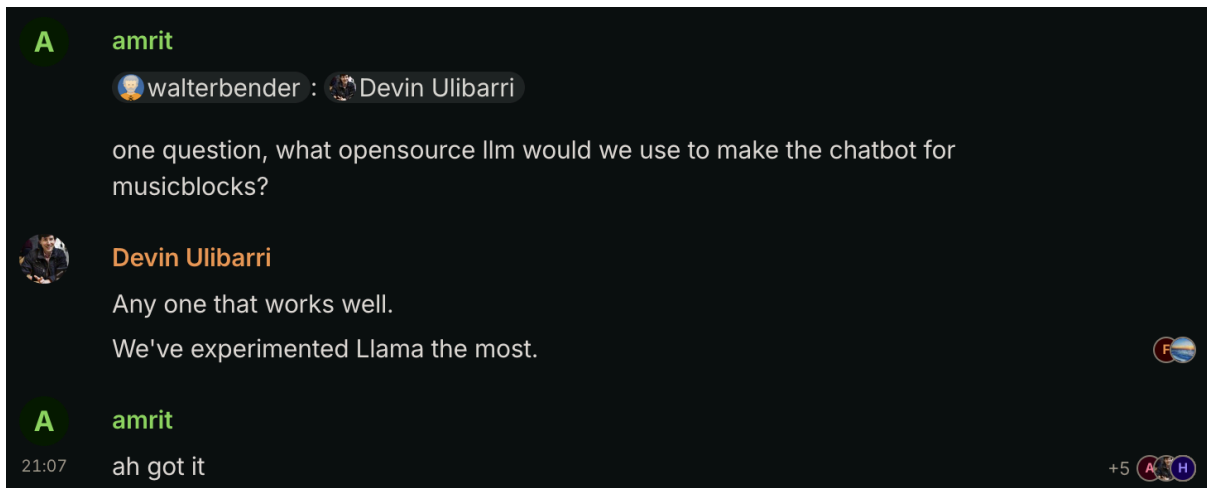
Scene AI

A SaaS I tried to build that uses a fine tuned open source model replicate to remove backgrounds from a video, with everything implemented from S3 storage, authentication and a payment method too. Let me know if you wish to try it out, I'll get you some free credits! The [link](#).

Many more great projects on my [github](#)!

[The project that I deeply wish to work on](#)

Interactive AI powered Chatbot and Debugger for Music Blocks



In an era where educational technology is rapidly transforming how young minds interact with and learn complex subjects, Music Blocks stands out as a dynamic platform that marries creativity with blocks to explore the fascinating world of music. This project envisions a groundbreaking enhancement: the integration of an AI-powered chatbot and a robust project debugger into the Music Blocks environment. By leveraging LLaMA (or any other open source model you guys recommend), an innovative open-source language model celebrated for its adaptability and high performance, our initiative seeks to empower users with immediate, context-aware assistance and detailed troubleshooting capabilities. Imagine a system where students, regardless of their technical proficiency, can not only receive intuitive guidance while navigating Music Blocks' diverse functionalities but also gain insightful, step-by-step debugging support that clarifies the intricate logic behind their projects. This synergy of creative exploration and technical refinement will not only lower the barrier to entry for newcomers but also offer seasoned users a richer, more engaging experience. Embracing the open-source philosophy of the Sugar Labs community, our proposal aims to set a new standard for interactive learning environments—one where technology and art converge seamlessly to foster a deeper understanding of music, coding, and digital creativity. It'll be a great chatbot. And a debugger.

1. Project Overview

Objective:

Enhance Music Blocks by integrating an AI-powered chatbot and a project debugger. These tools will:

- **Assist Beginners:** Provide real-time help, answer questions, and explain features.
- **Support Advanced Users:** Offer debugging insights to resolve issues in project logic and block connections.

Key Technologies:

- **Programming Language:** Python
- **Frameworks:** FastAPI for API endpoints
- **Platforms:** AWS for deployment
- **Core Idea:** Open-source LLMs, Retrieval-Augmented Generation (RAG), Chatbots, and Debugging Tools

2. Detailed Breakdown

A. Research and Analysis Phase

- **Understand Music Blocks:**
 - Study the existing codebase and typical project structures.
 - Identify common user errors and debugging challenges.
- **Review AI Chatbots and Debuggers:**
 - Analyze similar projects and approaches for integrating AI into creative platforms.
 - Determine the best practices to ensure usability and accuracy.

B. Architecture Design

- **System Architecture:**
 - **Modules:** Define separate components for the AI chatbot, project debugger, and API layer.

- **Data Flow:** Map out how data moves between Music Blocks, the LLM, and the front-end interface.
- **Model Selection & Customization:**
 - Choose an open-source LLM as the base model.
 - Outline plans to fine-tune the model on Music Blocks-specific data.
- **Integration Strategy:**
 - Use FastAPI to create RESTful endpoints for real-time interactions.
 - Design the interface to be easily integrated into the existing Music Blocks environment.
- **AWS Deployment:**
 - Plan for scalable deployment on AWS to handle inference requests.
 - Define monitoring and scaling strategies for the API.

C. LLM Training and Retrieval-Augmented Generation (RAG)

- **Data Collection:**
 - Gather Music Blocks project data, error logs, and user interactions.
 - Collaborate with the Sugar Labs community to obtain diverse use-case scenarios.
- **Training Strategy:**
 - Fine-tune the selected open-source LLM on the collected data.
 - Implement a RAG pipeline to enhance contextual understanding and mitigate hallucinations.
- **Accuracy Enhancements:**
 - Incorporate techniques (e.g., controlled generation, iterative feedback loops) to minimize hallucinations.
 - Set up evaluation metrics to continuously monitor model performance.

D. AI Chatbot Integration

- **Functionality:**
 - Answer user queries about Music Blocks features and common issues.
 - Offer creative suggestions to inspire new projects.

- **User Interface (UI):**
 - Design an intuitive UI/UX integrated within Music Blocks.
 - Ensure the chatbot is easily accessible and interactive.
- **Real-time Assistance:**
 - Implement dynamic response generation using the fine-tuned LLM.
 - Utilize RAG to fetch contextually relevant documentation and help resources.

E. Project Debugger Development

- **Error Identification:**
 - Develop algorithms to analyze project structures and block connections.
 - Detect common errors or misconfigurations in Music Blocks projects.
- **Debugging Assistance:**
 - Provide step-by-step suggestions to resolve detected issues.
 - Create visualization tools to help users understand debugging outputs.
- **Integration with Chatbot:**
 - Allow the chatbot to trigger debugging sessions.
 - Offer a seamless transition between chat-based guidance and detailed debugging information.

F. API Development with FastAPI (If not, we can proceed with a Node Js API too!)

- **Endpoint Design:**
 - Develop RESTful endpoints for chatbot interactions and debugger functions.
 - Ensure secure, efficient communication between the client (Music Blocks UI) and server.
- **Backend Integration:**
 - Integrate the LLM model and debugging logic within the FastAPI framework.
 - Prepare endpoints for model inference, error reporting, and feedback collection.

- **Performance Optimization:**
 - Optimize API responses to ensure minimal latency, especially for real-time interactions.
 - Implement caching and load balancing strategies where necessary.

G. Testing, Evaluation, and Iteration

- **Unit and Integration Testing:**
 - Develop comprehensive tests for each module (LLM, API, debugging tool).
 - Validate correct error detection and appropriate chatbot responses.
- **User Feedback:**
 - Run beta tests with community members.
 - Gather feedback and iterate on features to improve usability and accuracy.
- **Performance Metrics:**
 - Establish metrics for response time, debugging success rates, and overall user satisfaction.
 - Continuously monitor and adjust the system based on these metrics.

H. Documentation and Deployment

- **Documentation:**
 - Create developer and user guides detailing system architecture, API usage, and integration instructions.
 - Provide tutorials and examples to help new users and contributors.
- **AWS Deployment (If we use it):**
 - Finalize deployment on AWS, ensuring scalability, reliability, and cost-efficiency.
 - Set up monitoring dashboards and logs to track system performance.
- **Community Engagement:**
 - Share documentation and deployment strategies with the Sugar Labs community.
 - Encourage contributions and collaborative improvements.

3. Timeline and Milestones

1. Community Bonding Phase (Weeks 1–4):

- Engage with the Sugar Labs community and Music Blocks users.
- Refine project objectives and gather initial requirements.
- Set up a development environment and preliminary project planning.

2. Design Phase (Weeks 5–8):

- Finalize system architecture and module designs.
- Define API specifications and integration points.
- Prepare datasets for LLM training.

3. Development Phase (Weeks 9–20):

- **LLM Training & RAG Implementation:**
 - Fine-tune the model and set up the RAG pipeline.
- **Chatbot and Debugger Development:**
 - Build the core functionalities and integrate into Music Blocks.
- **FastAPI Endpoint Development:**
 - Develop and test API endpoints for real-time interaction.

4. Testing and Iteration (Weeks 21–24):

- Conduct unit and integration testing.
- Beta test with users and collect feedback.
- Optimize performance and address detected issues.

5. Final Deployment and Documentation (Weeks 25–26):

- Finalize documentation and user guides.
- Deploy the solution on AWS (if we are going to use it).
- Wrap up the project, present results, and prepare for potential future enhancements.

4. Risks and Mitigation Strategies

● Data Scarcity for LLM Training:

- *Mitigation:* Collaborate with the community to source diverse project examples and error logs.

● Model Hallucinations and Inaccuracies:

- *Mitigation:* Use RAG and iterative fine-tuning; incorporate user feedback for continuous improvement.

● Integration Challenges with Existing Music Blocks Codebase:

- *Mitigation:* Maintain active communication with core developers; implement incremental integration and testing.

- **Scalability Issues in API Deployment:**
 - *Mitigation:* Leverage AWS services for auto-scaling and performance monitoring.

5. Expected Outcomes

- **Enhanced User Experience:**
 - A robust AI chatbot that provides real-time assistance and creative suggestions.
- **Efficient Debugging:**
 - An integrated debugger that simplifies error resolution in Music Blocks projects.
- **Increased Accessibility:**
 - Lower barriers for beginners and a more streamlined workflow for advanced users.
- **Community Contribution:**
 - Open-source contributions back to the Sugar Labs and Music Blocks projects, to fix bugs and other issues.

Deliverables

This project is designed to deliver a comprehensive, state-of-the-art enhancement to the Music Blocks platform by integrating a finely tuned LLaMA-based AI, an intelligent chatbot, and an advanced project debugger. The deliverables span from the granular technical components to high-level community resources, ensuring that every aspect of the system is robust, scalable, and tailored for both novice and advanced users.

1. LLaMA Fine-Tuned Model for Music Blocks

Outcome:

A customized version of the LLaMA open-source language model, fine-tuned specifically on Music Blocks data, capable of understanding and generating context-aware responses to project-related queries and debugging scenarios.

Implementation Details:

- **Data Aggregation & Preprocessing:**
 - **Source Datasets:** Collect a wide variety of Music Blocks project code, user interaction logs, and error reports from the Sugar Labs community.
 - **Training Setup:** Utilize state-of-the-art frameworks (e.g., Hugging Face Transformers) to fine-tune LLaMA. Experiment with hyperparameters such as learning rate schedules, batch sizes, and gradient accumulation steps to optimize performance.
 - **Retrieval-Augmented Generation (RAG):** Integrate a RAG pipeline by indexing a curated knowledge base. This enhances the model's ability to retrieve accurate and contextually relevant data during real-time interactions.

2. AI-Powered Chatbot Integration

Outcome:

A dynamic, interactive chatbot embedded within Music Blocks, providing users with real-time assistance, feature explanations, and creative guidance.

Implementation Details:

- **Chatbot Architecture:**
 - **Core Engine:** Leverage the fine-tuned LLaMA model as the backbone for natural language understanding and generation.

- **Dialogue Management:** Design a conversation flow system that can handle multi-turn interactions, maintaining context across sessions.
- **Context-Aware Responses:** Utilize the RAG framework to fetch context-specific documentation, ensuring that responses are not only accurate but also actionable.
- **User Experience (UX):**
 - **Interface Design:** Develop a seamless and intuitive UI module within Music Blocks, using modern web frameworks. Prioritize responsiveness and ease-of-use with adaptive design for various devices. I am great with frontends too, so this will be great as well.

3. Project Debugger

Outcome:

An integrated debugger that can automatically analyze Music Blocks projects, identify syntax and logic errors, and provide step-by-step troubleshooting guidance. This will be a tough one.

Implementation Details

- **Error Detection Algorithms:**
 - **Static Analysis:** Develop static code analysis tools tailored for Music Blocks block-based projects. Identify common pitfalls such as misconfigured block connections and inconsistent variable use.
 - **Dynamic Analysis:** Simulate project execution in a sandbox environment to detect runtime errors, enabling the system to capture both compile-time and execution issues.
 - **Pattern Recognition:** Leverage machine learning techniques to recognize recurring error patterns, enhancing the model's predictive capabilities in identifying potential issues.
- **Debugger Interface:**
 - **Visual Feedback:** Create interactive visualizations that map out project structures, highlighting problematic areas and offering visual cues for resolution.
 - **Integration with Chatbot:** Enable the chatbot to initiate debugging sessions automatically when users describe issues, creating a seamless troubleshooting experience.

4. FastAPI Endpoints and Backend Integration

Outcome:

A robust set of RESTful API endpoints that allow seamless communication between the Music Blocks front-end, the LLaMA-powered chatbot, and the debugging module.

Implementation Details:

- **API Design and Development:**
 - **Endpoint Specification:** Design endpoints for key functionalities such as chat interactions, debugging queries, error logging, and user feedback.
 - **Error Handling & Logging:** Build comprehensive error handling routines and logging mechanisms to track API usage and identify bottlenecks.

5. Comprehensive Documentation and Community Resources

Outcome:

A complete set of technical documentation, user guides, and community resources that empower developers and users alike to understand, deploy, and extend the project.

Implementation Details:

- **Technical Documentation:**
 - **Architecture Blueprints:** Provide detailed diagrams and technical write-ups covering the system architecture, data flows, and module interactions.
 - **API Guides:** Develop comprehensive guides and code examples for integrating and interacting with the FastAPI endpoints.
 - **Training Manuals:** Document the LLaMA fine-tuning process, including data preparation, model training, evaluation, and deployment instructions.
- **FAQs & Troubleshooting Guides:** Develop a robust FAQ section addressing common user issues and providing troubleshooting tips if we find any along the way.
- **Open-Source Contribution Guidelines (If needed):**
 - **Contribution Roadmap:** Outline clear guidelines for external contributions, including code standards, pull request procedures, and issue tracking.

Final Thoughts

Each deliverable in this project is carefully designed to create an integrated, intelligent, and user-friendly system that elevates the Music Blocks experience. By combining LLaMA with robust debugging tools and a scalable API framework, we not only streamline the learning process but also empower users to explore and create with confidence. The comprehensive deployment ensures that this solution remains accessible, high-performing, while extensive documentation and community engagement foster an ecosystem of continuous improvement and collaboration. This ambitious yet pragmatic approach provides the edge needed to set a new standard in educational technology, delivering tangible benefits to both the Sugar Labs community and the broader world of interactive learning. Let's do this!

Impact on Sugar Labs

This project is designed to bring a range of tangible benefits to the Sugar Labs community and the Music Blocks platform. By adding an AI-powered chatbot and an integrated debugger, we aim to create a more interactive and supportive learning environment. Here's how it will make a difference:

- **Enhanced User Support:**
The chatbot will serve as an on-demand assistant that answers common questions and explains features in real time. For new users, this means less frustration when they get stuck, and for experienced users, quicker tips on how to resolve specific issues.
- **Improved Debugging Experience:**
Many users face challenges when troubleshooting their Music Blocks projects. The debugger will analyze project structures and identify errors, offering clear, step-by-step guidance on how to fix them. This makes the overall experience smoother and more rewarding.
- **Stronger Community Engagement:**
With tools that help users learn and solve problems faster, more people will feel comfortable experimenting and contributing. This increased participation can lead to more community-generated enhancements, fostering an environment where everyone learns from each other.

- **Boosting Educational Impact:**

Music Blocks is already a fun way to learn about music and coding. By integrating intelligent tools that simplify learning and troubleshooting, Sugar Labs can attract a broader audience—from complete beginners to seasoned coders—thus reinforcing its mission to provide quality educational experiences.

- **Encouraging Open-Source Contributions:**

The design of the chatbot and debugger is modular, meaning that once the project is up and running, other community members can easily add features or improve existing ones. This not only leads to a better product over time but also reinforces the open-source spirit of Sugar Labs.

Overall, the project is set to elevate the Music Blocks experience by reducing barriers to learning, speeding up problem resolution, and creating a vibrant, collaborative community around educational technology.

[Availability](#)

I can dedicate 35-45 hours per week to the project and will be active between Thursday and Sunday from 9 AM to 9 PM IST. I will have my exams in May-June period but that will be manageable where I will still dedicate 2-4 hours per day.

Conclusion

In conclusion, this project aims to make a significant impact on the Music Blocks platform by integrating an AI chatbot and an intelligent debugger. These tools are not just about adding cool features—they're about creating an environment where learning is made easier and more interactive. With immediate help available via the chatbot and a systematic approach to troubleshooting through the debugger, users can overcome challenges more efficiently and feel more empowered in their learning journey.

Moreover, by engaging deeply with the Sugar Labs community throughout the project, you're helping to create a culture of continuous improvement and collaboration. The modular nature of the solution means that once it's in place, it will pave the way for further enhancements from both new contributors and experienced developers alike.

The proposed timeline is designed to be realistic for a college student, balancing the workload with academic responsibilities while still ensuring the project is completed within a solid timeframe. Each phase—from initial planning and research to development, testing, and final deployment—has been carefully outlined to provide a clear roadmap to success.

Ultimately, this project not only enhances the user experience for Music Blocks but also reinforces Sugar Labs' commitment to open-source educational technology. It's a step forward in making learning more accessible, fun, and effective for everyone involved.

Being well versed in technologies and experienced in building projects, I'll ensure that we will build something we'd love to use.

Will be glad to be on board, thank you!