PromptSea

A Decentralized Prompt Marketplace for Generative AI

Pisuth Daengthongdee

Draft 1.0

Abstract

With the growing popularity of generative AI in content creation, millions of AI-works are being created everyday. Each work is made up of three core components, the textual description or prompt, the generative AI model, and the memory, although the memory component is often unused, but its potential significance in the future for the creation of more sophisticated works.

We propose a solution aimed at facilitating the generation of millions of commercially viable AI-works, aligning with copyright office regulations, and safeguarding the intellectual property of prompt creators. This is achieved through the conversion of the mentioned core components into a decentralized entity and using zero-knowledge technologies on zkSNARK and PLONK to streamline them all. This ensures a comprehensive balance, connectivity and benefits between the demand (prompt creators, end-users) and supply (original creators, operators).

Background

The advancement of large language models and generative AI has opened up a compelling opportunity, allowing anyone to produce high-quality creative works such as digital designs, graphic novels, animations and even music compositions.

For example, by utilizing a latent text-to-image diffusion model, we can produce photo-realistic images by providing text-based prompts that contain relevant keywords or key phrases associated with specific styles or elements of the desired image. The diffusion model then converts the text input into a latent space representation and progressively refines it through a series of steps, increasing image quality at each stage.

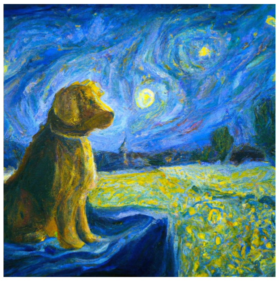

Using the prompt “dog, van Gogh, starry night, oil painting by James Gurney” in Stable Diffusion

And when combined with a large language model like GPT, we can achieve powerful capabilities in generating content, enhancing storytelling and facilitating content creation across multiple domains such as advertising, entertainment and graphic design.

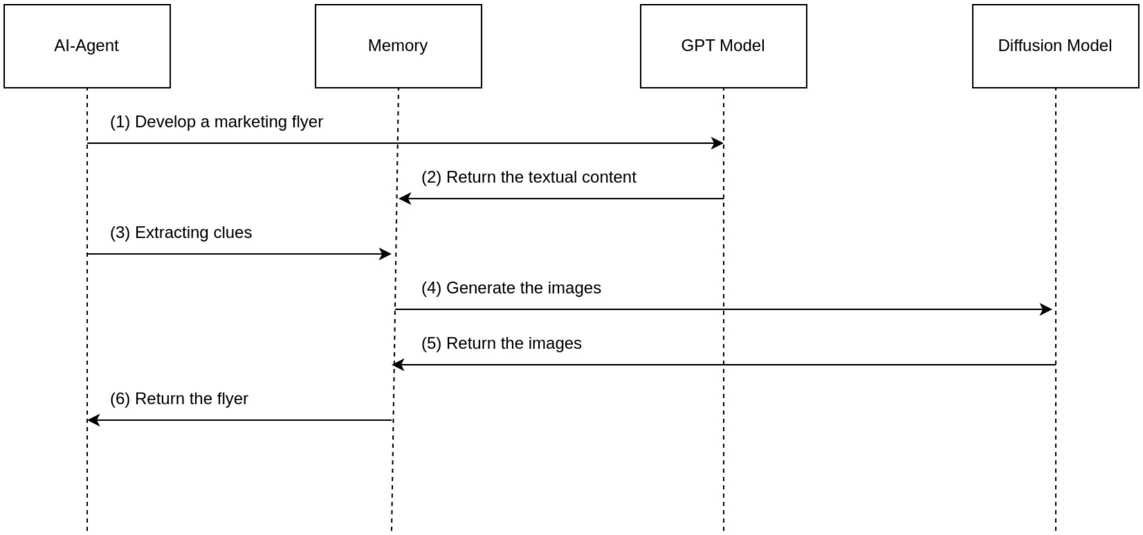

The basic flow demonstrates the utilization of multiple models to accomplish a target objective

When the technology is ready, we will witness the seamless integration in action, enabling the automation of the content generation process while maintaining high-quality standards. The AI-agent will play a crucial role in interacting with various AI-models and utilizing memory to store works in progress, ensuring efficient workflow and continuity.

However, the topic of AI-copyright and the commercial aspects have been subjects of ongoing debates across the industry and various sectors.

Prior to the initial draft of this whitepaper, PromptSea had been running for 6 months and was mentioned in some research papers. We gained valuable insights to identify key issues obstructing the advancement of the technology, which resulted in a large portion of AI-generated content being limited to leisure activities and encountering obstacles in terms of commercial utilization.

We can breakdown the major issues of each component as follows.

Prompt Stealing

This refers to the act of unauthorized copying or using someone else's text-to-image or text-to-text prompts or other AI-based content creation process without permission or proper attribution. Because creating a high-quality prompt that consists of a subject and several modifiers can be time-consuming and costly. This involves taking someone’s creative idea or description and using it as the basis for generating content through AI-models or systems.

There are a number of prompt marketplace operating in the market, including ours. It is possible for individuals to buy prompts from one marketplace and list them for sale on another marketplace. Such actions undermine the trust and credibility of prompt marketplace and diminish the value of AI-works overall.

Unauthorized Training Data

The training of AI-models often involves using large datasets that may include content created by individuals without their explicit authorization. This raises concerns regarding the ethical and legal aspects of using such data. While some text-to-image AI-models state their restrictions on commercial usage, there is still lack of assurance of the authenticity of these claims.

Moreover, the use of AI-models trained on unauthorized data sources raises concerns about the ethical implications of their output. Without proper authorization and consent, there is a risk of violating the rights and intellectual property of original creators.

Undoubtedly, AI-generated content is a disruptive technology, it is to be expected that some may resist or oppose it from individuals or groups who perceive potential losses of disadvantage it may bring. Within the discipline of change management, one approach to making change happen is by increasing the driving force that supports the change and reducing the resistance forces that impede it.

Consequently, we propose a solution that creates a win-win situation for both sides by addressing the needs and concerns of both parties involved, our solution aims to strike a balance and foster a relationship among them.

Overview

The fundamental purpose behind employing a decentralized model is to allow application developers to step aside, facilitating a self-regulating system where demand and supply can find equilibrium.

For instance, let’s consider a cryptocurrency exchange as an example. In a decentralized exchange, the responsibility of providing liquidity rests with the participants while end-users benefit from automatically calculated rates determined by the available liquidity within the system.

Requirements

To achieve the stated objective, the system must meet the following requirement:

- Verifiable privacy-preserved - Tradable prompts should not be disclosed to anyone while retaining the ability to execute on specific AI-models in a verifiable manner.
- Ethical AI model training - Conducting AI model training in a responsible manner that includes obtaining authorization and consent from the data owner.

- Maximizing prompt creation - Enabling individuals to assess prompts or input texts on privacy AI models and monetize them.
- AI-native memory - Enabling conversion of prompts and consent data into vector representations, preventing duplicate listings.

Security Model

In order to create a decentralized marketplace with the ability to develop privacy-focused AI-models, the system adopts a distinctive distributed model similar to Chainlink, where a community of providers collaboratively source and contribute to the ecosystem.

Due to the dynamic and diverse nature of the AI landscape, it is not feasible to have a unified network that collects and maintains the same data, comparable to the synchronization of full nodes in blockchain and rollup networks.

Architecture

As we implement a distinctive distributed model, off-chain nodes running by the community play an important role in the network architecture. These nodes located outside the main blockchain, perform crucial tasks such as indexing, data processing, embedding and storage.

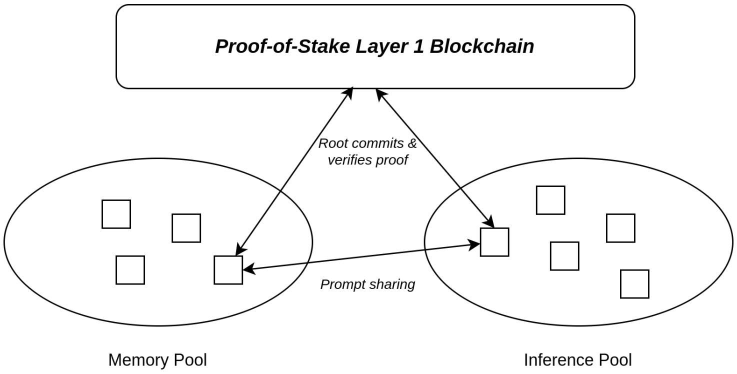
Proof-of-Stake Layer 1 Blockchain
Root commits & verifies proof
Prompt sharing
Memory Pool
Inference Pool

The architecture of PromptSea's decentralized network

All nodes aggregate into a pool called the syndicate, which consists of three distinct type of pools:

- Memory Pool - The pool operates as a vector database cluster, with all nodes running an eventually consistent document database software together with a vector library. When data is injected into the pool, one of the nodes performs text embeddings and then encoding operations using Circom's circuit for further verification through zero-knowledge proofs.
- Inference Pool - The pool is responsible for running services built from the model and supports running tasks on transformer, sentence-transformers and diffusers libraries from HuggingFace. Similar to the memory pool, it necessitates encoding all operations for on-chain verification and facilitates the distribution of fees to the pool, data owner and prompt owner.

- Utility Pool - The pool is utilized for various purposes and should be operated by trust entities rather than the community to handle sensitive tasks like generating asymmetric keys for end-users, managing keys and access control.

The Proof-of-Stake layer 1 blockchain serves as a robust and trustless infrastructure for on-chain verification, facilitating transparent and equitable payments among participants. As the project is currently in its early stages, we’re exploring multiple blockchains, including the Testnet. However, we will carefully select the one blockchain that offers the most advantages and benefits for our project.

zkSNARK and PLONK are employed to establish a trustless marketplace, ensuring secure and transparent interactions between users, prompt creators and data providers.

Actors

In the system, there are several key actors who play distinct roles and contribute to the overall functioning of the platform:

- Data providers - Individuals or entities who possess and own the data used to train the generative model such as artworks, illustrations and game assets. All contributions are tagged and verifiable in terms of ownership, allowing data owners to receive payments corresponding to their contribution’s portion when the model is executed.
- Prompt creators - They actively explore the potential of AI models and create prompts that showcase their capabilities and can list them for sale. Similar prompts that produce the same output are prevented from being listed.
- Node operators - They are responsible for running the necessary infrastructure and maintaining the resources needed to support the execution of tasks and operations on memory, inference and utility pools and contribute to the overall functioning of the decentralized network.

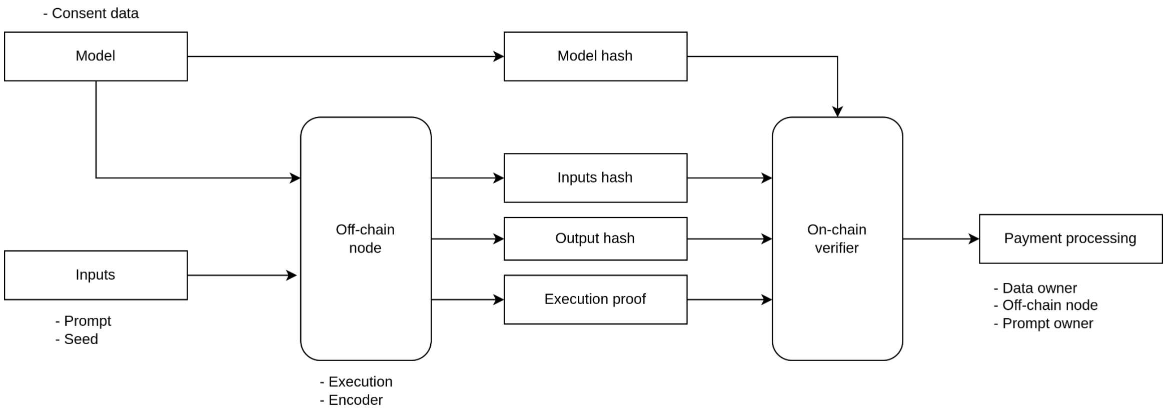

All actors work together in a verifiable manner through the use of zkSNARK and PLONK

The end-users can then browse AI-models and choose prompts that can generate images or specific outputs without knowing the exact prompt. Some models may not require prompts, and in such cases, there will be direct interaction between the user (via the operator node) and the AI model.

Memories

Often used with vector databases, unlike traditional databases that rely on human-language to query the data, a vector database allows the conversion of data into a vector representation and enables querying by AI.

Some AI-models, such as the diffusion model, have a built-in vector store, while external vector databases can be used when interacting with multiple models at the same time. The use of a vector database can effectively prevent prompt duplication and ensure efficient utilization of prompts.

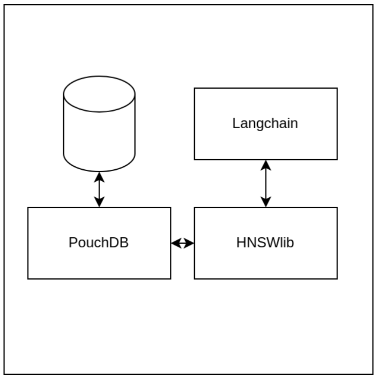

The memory node incorporates PouchDB, Hnswlib and Langchain

In the system, the memory pool is utilized instead of individual memory nodes. The memory pool is a cluster of vector databases designed for efficient storage and retrieval of vector representations in a Web3 fashion, ensuring access to vector data and facilitating various AI-driven operations across different pools.

Each memory node incorporates PouchDB, Hnswlib and Langchain to create an efficient and reliable memory system for the entire system. It has low memory consumption and small footprint compared to Chroma, another vector database solution for non-Web3 that incorporates Clickhouse that requires more resources.

The reasons for choosing these components are provided below:

- PouchDB - It operates as a compact non-relational database with eventual consistency, it offers a small footprint and has a crypto-pouch plugin for document encryption using the xsalsa20-poly1305 algorithm, ensuring the data is fully protected.
- Hnswlib - An open-source library implementing the HNSW algorithm for approximate nearest neighbor search, widely used in image retrieval and natural language processing.
- Langchain - A framework for developing with large language models offering modules for document parsing, model interaction and standard interfaces.

The memory node also provides a specific endpoint for Langchain-formatted vector data that allows developers to directly access it as a remote source within their applications.

Applications

There are various applications with PromptSea's decentralized network system, serving different purposes and catering to diverse user needs.

Commercially-viable AI-works

Commercially-viable AI-works are made possible through authorized access granted from the original intellectual property owners. This allows for the utillzation of AI models in a legally compliant and ethically sound manner, enabling businesses and individuals to confidently leverate these models.

As a result, the commercial-viability of AI-works provides opportunities for various industries such as art, gaming, content generation and many others. They can leverage AI models to enhance their processes, create unique experiences and drive innovations.

Generative QA for Specific Niches

The training data for GPT models primarily consists of public information on the internet. In contrast, businesses have the opportunity to train their own AI models using proprietary data such as patents and archival material.

For instance, television stations often possess extensive archives containing footage capturing moments from various time periods. By importing scripts and relevant content into the memory pool, users can ask questions and receive accurate responses based on the right and specific knowledge contained within those archives.

Roadmap

We have divided the development process for a decentralized marketplace of self-sovereign AI models into 4 phases. The first phase has been successfully accomplished and we are in the middle of phase two at the moment, marking a significant milestone in our journey.

Phase 1	Phase 2	Phase 3	Phase 4
Initial launch and support prompts for external models	Implementation and deployment of the Memory Pool	Development and integration of the Inference Pool	Building and integration of ZKML-based models
(Completed)	*(In progress)*	*(Q3/23)*	*(Q4/23)*

The token generation event will occur after Phase 4

Awards and Recognition

- PromptSea has been incubated by FuelArts x Tezos accelerator from February until April 2023
- The memory node has won one of grand prizes at the FVM Dataverse Hack on May 2023

References

https://www.researchgate.net/publication/368664479_Prompt_Stealing_Attacks_Against_Text-to-Image_Generation_Models
https://worldcoin.org/blog/engineering/intro-to-zkml
https://research.chain.link/whitepaper-v2.pdf
https://medium.com/@danieldkang/trustless-verification-of-machine-learning-6f648fd8ba88
https://medium.com/@danieldkang/open-sourcing-zkml-trustless-machine-learning-for-all-f5ee1dbf2499
https://github.com/timoftime/zk-SQL
https://arstechnica.com/information-technology/2023/02/us-copyright-office-withdraws-copyright-for-ai-generated-comic-artwork/