Rumi Preseed Investment Memo (March 2025)
Rumi is indexing all the world's videos to enable real-time AI-powered content experiences.

Investment Overview

- **Rumi is a decentralized video indexing network**. Users (miners) run Rumi's browser extension while watching videos on streaming platforms like Netflix or Youtube. The extension uses edge ML to produce context-rich text annotations. These annotations - not the videos themselves - are sent back, stored and processed in Rumi's database, where they are made available to AI agents seeking to deliver highly-contextual experiences / services / advertisements to end-users.

- **Bootstrapping consumer demand with a *universal remote ("UR")* mobile app.** Rumi is bootstrapping demand for the network by launching the first AI-powered UR app. UR apps allow users to control any smart TV from their mobile phone, eliminating the need for a physical remote. These apps see 4m monthly downloads globally at low user acquisition costs; most users search the app store for them after losing their physical remote or when it runs out of battery. Rumi's app will be the first AI-powered UR app, capable of acting on voice prompts from users with context.

- **Proprietary audio fingerprinting technology enabling real-time agent-driven experiences.** Rumi's team has spent the past two years, prior to incorporating the company, developing audio fingerprinting technology that is able to identify what a user is watching and how far into the video they are with very low latency (<1s) and high accuracy (>98%) with only a 400ms audio clip in ideal conditions. This technology enables experiences whereby agents are able to tap into a rich database of contextual information corresponding to exactly what users are watching in real-time.

- **Long-term path to become the decentralized Nielsen.** Nielsen generates $1.5B+ in annual EBITDA by publishing industry-standard benchmarks for ad reach and effectiveness. Both the buy-side (brands & agencies) and sell-side (content owners & distributors) use Nielsen's data as the basis for pricing advertising slots. Nielsen's data is based on a network of ~50k homes that earn ~$15/mo for installing a small device that monitors and reports TV activity. Using DePIN scaling playbooks, we believe Rumi can reach millions of homes and achieve better, more statistically-significant benchmarks than Nielsen, ultimately creating 3-sided network effects between ad buyers (brands), ad sellers (content distributors), and AI agent developers.

- **Serial entrepreneurs who understand the customer problem inside & out.** Co-founders Niko Cunningham and Matt Poreda have been building at the cutting edge of crypto, AI and consumer for the better part of a decade. Niko was most recently the senior director of product at Udacity, where he led all AI and consumer bets from '16-'19 driving growth from 10k to 100k students. Before that, he founded one startup to successful M&A exit (Nerd Swagger: digital scrapbooks) and led product at another successful M&A exit (Koemei: annotations for university lectures). Matt was the first employee and COO at Omni Network, a consumer-focused chain abstraction protocol, where he led efforts to launch $OMNI which achieved a Binance listing at a $1B FDV.

- **EV3 is co-leading Rumi's $4.7m pre-seed round** alongside a16z CSX. The financing will provide the company runway to build and scale their supply-side indexing capabilities, develop the (first-party) AI-powered universal remote mobile app, and attract (third-party) developers to create agentic applications on top of real-time streaming and TV content powered by Rumi.

Network Overview

Rumi's network is comprised of three parts:

1. **Supply-side miners:** that download Rumi's browser extension and earn rewards for contributing to the network's database of context-rich video annotations.
    - Rumi's database was initialized with an existing corpus of scripts, closed captions, translations, and other metadata of video content. When a miner watches a stream with Rumi's extension running, Rumi's annotations database expands in reach (via indexing new videos) and in depth (by adding more context to existing videos). This creates a compounding moat since future competitors will need to start catalogues from scratch.

2. **AI agent developers:** that leverage Rumi's real-time fingerprinting (to identify what users are watching) and annotations database (to learn more context about what users are watching).
    - The first agentic app on the network will be developed by Rumi itself, i.e. the universal remote app described above. The app will enable users to control any smart TV using natural language, e.g. using queries like "rewind to the last first down", "what other movies has this villain been in?", "are there any spoilers in the next 10 minutes?". Users will earn inflationary token rewards for using the app, bootstrapping early demand.

3. **Demand-side consumers:** that use agentic apps and experiences powered by Rumi.
    - Apps powered by Rumi leverage privacy-preserving edge ML to identify what users are watching without jeopardizing their broader privacy. In testing, Rumi's models are able to fingerprint content from TVs up to 15 meters away at the lowest volume level, while isolating away user conversations and other background noise. By running ML models locally on users' smartphones, Rumi ensures that only TV/video annotations are shared.

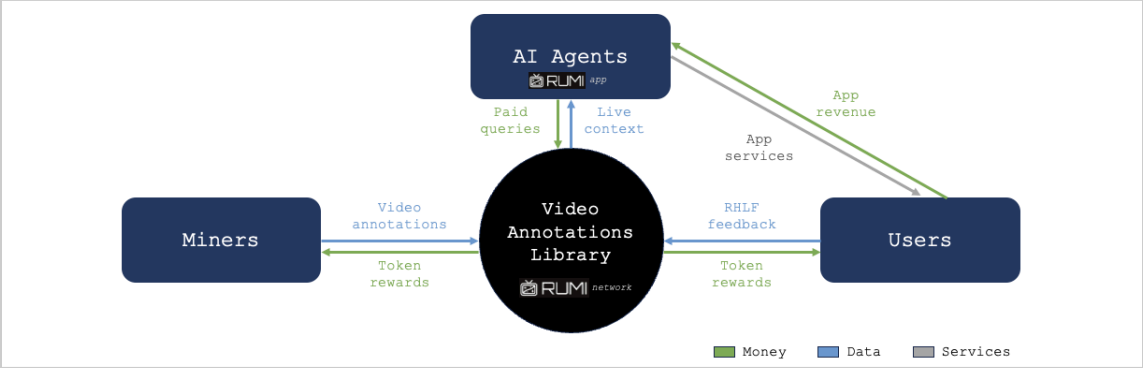

Rumi has the potential to create a 3-sided flywheel of miners (indexers), developers (agents), and users that aggregates the most expansive and context-rich database of video annotations on earth:

- As the database expands in breadth and depth, Rumi will support more powerful agentic experiences than competitors with smaller catalogues and surface-level annotations, ultimately driving higher returns/ROAs for agent developers and attracting further development.

- As users submit organic queries to agentic apps built on Rumi, the network "fetches" any missing context from miners and appends it to the database. This enriches the database with even more context for future users, driving an RHLF-like flywheel of defensibility.

WTNTB

We are underwriting the following for Rumi network on a 6-10 year time horizon:

- 250k agents/developers each paying $175 per month for access to Rumi's contextual database
- This implies a few hundred users per agent and a few cents per-user per-day in fees paid to Rumi
- 15x net revenue multiple: Nielsen was acquired for 11x EBITDA - the equivalent metric - in 2022

Rumi Network	2026-2027 Series A	2028-2029 Series B / TGE	2031-2032 Return of Capital	2034-2035 Full Exit	
Paying agents	10,000	50,000	250,000	1,250,000	<< assumes 400% growth at every stage
(x) Users per agent	100	200	400	800	<< assumes 100% growth at every stage
Active users	1,000,000	10,000,000	100,000,000	1,000,000,000	
(x) Monthly calls per user	30	45	68	101	<< 50% growth at every stage
(x) $ cost per call	$0.010	$0.008	$0.006	$0.004	<< assumes 25% decline at every stage
Network ARR ($m)	$4	$41	$456	$5,126	
Assumed multiple	25x	20x	15x	10x	<< lower growth premium at larger scale
Implied FDV ($m)	**$90**	**$810**	**$6,834**	**$51,258**	
Implied fund MOIC	0.1x	1.3x	11x	82x	<< assumes 8% ownership and $50m fund size
% of Nielsen '22 EBITDA	*0%*	*3%*	*30%*	*340%*	*<< Nielsen generated $1.5B RR EBITDA in Q3'22*
% of Nielsen '22 EV	*1%*	*5%*	*43%*	*320%*	*<< Nielsen was bought for 11x EBITDA by PE in Q3'22*

These scenarios imply the potential to return 10x+ the fund in the upside case, surpassing Nielsen.

We believe Rumi-the-network has a stronger moat and can be a far bigger business than Rumi-the-app. Based on napkin math, Rumi app could return the fund in the long-term bull / upside case but has a far lower ceiling than the network. The goal of the app is to bootstrap demand on and showcase capabilities of the network in the early days. We are underwriting 5m paid users at a $50m ARPU and 3x revenue multiple in the bull case, which is roughly 2-4x the total size of the UR market today by revenues.

Rumi App	Base	Upside	Bull	
Active users	5,000,000	20,000,000	100,000,000	<< non-AI universal remote (UR) apps see ~75m annual downloads
(x) % Paid	20%	10%	5%	<< current UR apps see 5-10% paid (lots of dark patterns)
Paid users	1,000,000	2,000,000	5,000,000	
(x) ARPU	$50	$50	$50	<< current UR apps charge $50-150/yr
App ARR ($m)	$50	$100	$250	
Assumed multiple	5x	4x	3x	<< assumes lower growth premium at larger scale
Implied EV ($m)	**$250**	**$400**	**$750**	
Implied fund MOIC	0.4x	0.6x	1.2x	<< assumes 8% ownership and $50m fund size

Most scaled companies in the advertising data and analytics space have been acquired by private equity firms. In particular, from 2019-2022 syndicates led by Elliot Management, Hellman & Friedman, and BainCapital acquired $25B of advertising data businesses (with leverage) at 6-11x EBITDA multiples.

	Nielsen	IRI	Kantar	Pitchbook
Sector focus	Media & Advertising	Retail & POS	Brand & Consumer	Private markets
PE acquirer	ELLIOTT	H&F	BainCapital	MORNINGSTAR
Acquisition year	2022	2022	2019	2016
Annual revenue	$3.5B	n/d	$3.5B	$550m
Purchase Price	**$16B**	**$5B**	**$4B**	**$3B**(1)

Note: Pitchbook was acquired for $225m in 2016; $3B assumes 5x multiple of latest revenues.

Key Risks

1. **Indexing works for 'persistent' video content, not for 'disposable' AI-generated videos**. Assuming a growing share of video is generated "on the fly" by genAI models, indexing becomes less valuable as there is no persistent transcript or annotations for agents to ingest. However, we believe the vast majority of video will continue to be persistent in nature. TVs still represent the majority of screentime in the US - even more than mobile - and only a very small minority of TVs are equipped with chips anywhere near powerful enough to generate real-time video.

2. **Bootstrapping demand on the network.** Rumi has a two-pronged strategy to bootstrap demand through: 1) their first-party universal remote mobile app, and 2) integrations with third-party agents infrastructure like Virtuals and ai16z. We believe Rumi will be able to quickly win market share in the universal remote market given most competitors are bootstrapped apps built by agencies, full of dark patterns for monetization and engagement. Rumi will be the first AI-powered UR app, and the first to offer token rewards for users, in addition to having a best-in-class UI. Universal remote mobile apps see 4m+ monthly downloads globally with little marketing spend.

3. **Navigating future conflicts with content owners.** From the perspective of a content owner like Netflix, Rumi is a net positive since users are more engaged and churn less when there's an entire ecosystem of agentic apps built on top of their content library. On the other hand, Netflix will see 0% of the advertising revenues (i.e., agent fees) generated by the Rumi network. We believe this will incentivize content distributors to ignore Rumi at the beginning, and eventually partner with Rumi for a revenue-share as soon as the network is seeing meaningful activity.

4. **Overcoming Nielsen's direct integrations moat** (e.g., DIRECTV, DISH, Vizio, Roku). Nielsen has strategic partnerships with several of the largest players across TV, streaming, and cable industries that gives them access to proprietary data and distribution channels for their devices. Rumi will need to overcome this by 1) building a highly performant architecture that integrates seamlessly across TV hardware & software without compromising the user experience, and 2) using DePIN playbooks to rapidly scale beyond the level that's feasible for Nielsen (~50k homes).

5. **Competition from unexpected places** (e.g. Amazon). Rumi's audio fingerprinting technology provides a temporary technical advantage in the medium-term but will eventually be replicated by competitors, if it hasn't been already. Amazon, for example, has an install base of 100m+ smart speakers capable of running similar audio fingerprinting models, and is aggressively muscling its way into other parts of the advertising value chain. Other companies like Roku, which controls ~40% of the smart TV OS market, are also well-positioned to compete against Rumi but face the innovator's dilemma of cannibalizing their existing direct ad sales businesses.

6. **Privacy concerns and regulatory risk** are both relatively limited. Rumi's indexers extract audio directly from soundcards, removing any ability to snoop on user conversations. For audio finger-printing, Rumi's edge ML models compare a hash of a single 400ms recording against the global content database. Because the clips are short (400ms) and only a hash is compared, there's very limited potential for meaningful information leakage. From a regulatory perspective, the main risk is being sued by content owners; however, the transformative use doctrine of the fair use defense states that content can be re-used for commercial purposes so long as the indexed content is a sufficiently large "transformation" from its original form. Rumi's indexers run ML models on their home computers at the edge and only send the context-rich text annotations back to the network.