

WHITE PAPER: THE OPPORTUNITY IN ACADEMIC AUDIO

Knowledgine, Inc.



Christian Balevski – CEO
Matthew Rodriguez – CTO

Abstract

Knowledgine is a modern university knowledge management platform that solves real problems for universities, professors, students, and learners everywhere. Knowledgine recognizes that many universities have benefited from the OpenCourseWare (OCW) movement, but also understands the extent of resources involved in OCW. Accordingly, Knowledgine has two main goals in focus:

- Drastically reduce the cost, resources, and time required to create high quality, globally accessible content.
- Deliver academic knowledge to students and learners in a modern and intuitive way.

Why: Our core philosophy is that great professors exist at every university. Education technology should be the means by which academic leadership is democratized, not consolidated because of prohibitive costs.

How: In response to market data and trends, Knowledgine takes a fundamentally different approach to the way in which it captures and distributes information to learners. By focusing on professor speech, Knowledgine delivers *sound* education through a process of automation and machine learning. This approach simplifies the work required by professors, and offers universities a cost-effective and scalable way of achieving local, regional, and global objectives.

Contents

Introduction	4
Opportunity in Audio	5
Knowledgine Platform – Overview	5
Benefits of Knowledgine	5
Platform Technology Stack	6
Lecture Capture	6
Professor Studio	6
Lecture Recording and Uploading	6
Lecture Transcription	7
Information Organization	8
Topics and Key Phrases	8
Tags	8
Playlists	9
Profiles	9
Systems in Development or Testing	10
Content Distribution and User Experience	10
Web Application	11
Search	12
Cloud Hosting	13
Product Roadmap	13
Professor Intellectual Property	13
Business Model	14
Conclusion	14
References	15

Introduction

Everyday, professors across the world produce high-quality information consisting of novel perspectives, transformative opinions and meaningful explanations. Only select universities have the resources to capture this content and benefit from its distribution.

For example, Massachusetts Institute of Technology (MIT), one of the biggest champions of OCW, have this to say about it upon reflection¹:

Over the past 10 years, OCW has moved from a bold experiment to an integral part of MIT. Currently, more than **93% of undergraduates and 82% of graduate** students say they use the site as a supplement to their course material or to study beyond their formal coursework. **Eighty-four percent of faculty members** use the site for advising, course materials creation, and personal learning. **More than half of MIT alumni** report using the site as well, keeping up with developments in their field, revisiting the materials of favorite professors, and exploring new topics. Open publication of course materials has become an ordinary element of scholarly activity for MIT faculty, and the ubiquitous availability of that curriculum to our own community has become the everyday reality of teaching and learning at MIT.

Universities with significant funding or endowment have been able to cover the cost associated with producing, capturing, hosting, and distributing their course content. The total annual cost of MIT OpenCourseWare is about \$2.7 million². Due to the mounting financial concerns of OCW, MIT, in conjunction with Harvard University, created EdX, a nonprofit Massive Open Online Course (MOOC) platform designed to embody the core principles initially pioneered by the leaders at MIT. Although the general success and impact of MOOCs is well accepted, some argue against their inclusivity due to the prohibitive costs of production and their general effectiveness in student retention and course completion.

A recent study done by Fiona M. Hollands and Devayani Tirhali of Columbia and Brown University, conclude the average cost to produce a MOOC is between \$203,770 - \$325,330³. The cost of production is compounded through the use of design and support teams, management, and expensive equipment.

Completion statistics of MOOCs don't bode well for their cost justification either. Some MOOCs have dropout rates as high as 97%⁴. Understandably, a 3% completion rate of 12,000 students is still very impressive. However, Knowledge believes this data points to the desire for many individuals to discover and explore content rather than committing to and completing courses. Knowledge embodies this philosophy as a core design component of its platform, catering to the needs and desires of global learners rather than forcing antiquated methods.

Given these costs, many universities will gravitate towards Learning Management Systems (LMS) and Content Management Systems (CMS) such as MediaSpace, Blackboard, and SonicFoundry to share their lectures. While these systems are comprehensive in their feature set, their outdated user experience and siloed nature often result in minimal view counts and unapproachable information. Ergo, this trade off between cost and quality is self-evident. *QED*

Whether producing online courses or using LMS or CMS, the fundamental problem for universities and professors is still the same: **how do we effectively deliver the right information to the right audiences.**

Opportunity in Audio

A brief walk around campus will quickly illustrate the ubiquity of audio consumption by students. According to a 2018 Nielsen study, Americans listen to 32 hours of music per week⁵. This doesn't include the 5.1 hours a week⁶ podcast listeners spend listening to podcasts or audiobooks. Voice is often considered to be the next big thing by industry experts and analysts alike. Saint Louis University in Missouri recently announced it will be putting Amazon Alexa enabled devices in every student living space on campus⁷ (over 2,300 smart speakers). This decision by a higher education institution reflects the market opportunity in voice and spoken content. Universities that continue to innovate in this regard will pioneer the new form of global education, empowering learners and professors in the process.

Knowledgine Platform – Overview

The Knowledgine platform is a full end-to-end pipeline that solves the problems around lecture capture, information organization, and content distribution. By using advanced technologies in automation, machine learning (ML), and crowdsourcing, **Knowledgine can ingest raw, unstructured professor speech and deliver ready-to-consume packets of knowledge to learners globally.**

Knowledgine's platform consists of 3 core areas:

- **Lecture Capture:** Professors can record or upload any lecture within Knowledgine's web application.
- **Information Organization:** Knowledgine uses a refinery inspired process for data organization where each level of refinement results in more structured pieces of consumable knowledge.
- **Content Distribution:** Knowledgine's open access proprietary web application delivers content to users on campus or across the country.

Benefits of Knowledgine

- **Simple:** One-click record, one-click listen.
- **Cost-effective:** Free for universities and professors.
- **Scalable:** Cloud hosted and able to meet the growing needs of your university.
- **Accessible:** Transcripts, metadata, future multi-language support.
- **Open:** Contribute to the betterment of society and global economy.
- **Cross-disciplinary:** Facilitate learning across disciplines to inspire creative thought.
- **Organized:** Topics, tags, playlists, and profiles.
- **Flexible:** Works alongside LMS and CMS systems.
- **Persistent:** Lectures are stored and backed up in the cloud.
- **Secure:** Encrypted data storage and transmission.

Platform Technology Stack

Lecture Capture

Knowledgine recognizes that professors do a lot of important work, and often don't have time to learn and explore new technologies. With that in mind, Knowledgine's lecture capture system is designed to be simple and intuitive, focusing on the fundamental part of lecture capture: speech.

- **Professor Studio:** A centralized place to create, access, and manage courses and lectures.
- **Lecture Recording and Uploading:** A professor can upload an existing recorded lecture or record a new lecture with Knowledgine's web based recorder. Lecture audio is securely uploaded to the cloud where it is processed and stored.
- **Lecture Transcription:** Knowledgine recognizes that **accessible content is better content**. This is why Knowledgine has automated the process of lecture transcription, with competitive results in terms of accuracy, time, and cost.

Professor Studio

Knowledgine's Professor Studio is an exclusive web service only available to university professors and academics. It is a centralized place where professors can create and manage courses, share resources through course links and descriptions, and upload and record lectures.

Knowledgine internally validates professors before they are given access to the Professor Studio. This insures that only legitimate professors from accredited universities share lectures on Knowledgine.

Lecture Recording and Uploading

Lecture capture is Knowledgine's core feature for professors. It is designed to be simple and requires limited resources. Whether in the classroom or in the office, the Professor Studio offers everything a professor needs to record a lecture, podcast, or clip.

Within the Professor Studio, a professor may simply create a new lecture by selecting the associated course, choosing a recording device (whether internal or connected), and then clicking record. This feature is designed to work with existing classroom equipment that is connected to the recording computer, and also with the built-in microphone of many modern laptops.

Professors also have the ability to pause and resume their recording, giving them flexibility and control if they want to discuss something confidential or if there is a break in class, for example. Furthermore, professors may add **tags** — labels associated with specific segments within the lecture (detailed in the *Information Organization* section of this white paper) — to their lectures.

Recorded lectures are available for both upload and download. Professors may optionally download the lecture if they want to edit it with third-party software and upload it later. Lectures uploaded to the cloud are securely processed, organized, and made available for discovery within the Knowledgine ecosystem.

Lecture Transcription

Knowledge utilizes IBM Watson's cloud based transcription service to deliver fast and accurate speech transcription. With this service, error rates are as low as 5.5% (human parity is 5.1%)⁸ and it continues to improve as more data is added to the system. Alone, IBM's cloud solution offers high accuracy but processes speech in roughly real-time — i.e. a 90-minute lecture will take 90 minutes to transcribe. Understanding the importance of fast access to accessible content, Knowledge developed pre- and post-processing algorithms that improve the transcription time by up to a factor of 30 — i.e. a 90-minute lecture will take 3 minutes to transcribe. This novel solution provides high accuracy transcripts with near instant availability.

Transcripts are also fully interactive and searchable, allowing learners to jump to relevant parts of a lecture quickly.

Utilizing ML systems for cloud based transcription, cost of service can be reduced by 100X: \$1.20/hour vs \$2/minute with traditional, human based, transcription services. Knowledge is actively researching methods to further reduce the cost and improve content accessibility.

The screenshot displays a user interface for a lecture transcript. On the left, a course overview for 'Virology - BIO302' by Vincent Racaniello is shown, including the university (Columbia University), dates (Jul 5, 2017 - Jul 12, 2017), and time (8:00 AM - 9:00 AM, Mon Wed). A list of lecture topics is provided, with 'What is a virus?' selected. The central area features a video player with a transcript overlay. The transcript is divided into segments with timestamps and text. The right sidebar contains a 'What is a virus?' section with various interactive options like 'Viruses are everywhere', 'Goals for this course', and 'Number or Viral Particles'. At the bottom, there is a 'Create Tag' section with fields for Name, Start Time, and End Time.

FIGURE 1: DISPLAY OF INTERACTIVE TRANSCRIPT ON LECTURE

Information Organization

The foundation of Knowledgine’s information organization strategy begins with transcription. Once text has been generated for a lecture, we apply a variety of natural language processing (NLP) methods and machine learning to understand lecture context and intent. Knowledgine couples this strategy with the power of crowdsourcing to provide high-quality clips of knowledge that are easily consumable and searchable. **All of this is done without any additional work by professors or staff.**

- **Topic and Key Phrase Extraction:** Word embedding and methods in NLP are used to develop a concise list of important topics and phrases covered in each lecture.
- **Tags:** (Patent-Pending) User generated labels attached to specific segments within lectures to make information discoverable.
- **Playlists:** (Patent-Pending) User curated lists of tags that allow individuals to express their own thoughts and ideas through the lens of academic understanding.
- **Profiles:** Collections of tags, playlists, lectures, and courses all organized by the owner of the profile.
- **Systems in Development or Testing:** Entity Recognition, Student Name Filter, Sentiment Analysis

Topics and Key Phrases

Once transcribed, subsequent lecture processing analyzes the text to identify keywords and topics. This system uses a combination of supervised and unsupervised machine learning, as well as proprietary statistical models to deliver a comprehensive list of topics and keywords discussed throughout the lecture. For instance, processing a lecture titled “Digital Technology and the Future of Work” extracts key topics like *automation*, *technological unemployment*, *globalization*, *big data*, and *accountants*. Worth noting, it is important to keep topic lists concise, otherwise every lecture begins to associate with every topic, obfuscating the foundational knowledge within.

Tags

Tags are the core feature that allow users to contribute to the organization and identification of valuable information within lectures. Each tag is comprised of a title, a start time, an end time, and a parent lecture. Upon tag creation, whether it be by student or professor, other users will be able to access the tag. When a tag is clicked, the lecture begins playing at the specified start time, and will optionally skip to the next tag once the end time is reached.

Although this system can be supplemented through the use of automation technology, Knowledgine believes in the importance of user contribution. Tags are not just indexed words, topics, or definitions but they are an expression of creative freedom for our users. Users index pieces of information that are useful to themselves, with labels that make the most sense to themselves. Consequently, the same piece of information may be redundantly tagged under different labels. It's not just a system for finding definitions; it's a system with human contribution, for human interaction and connection.

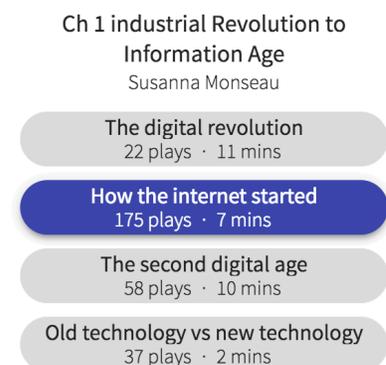


FIGURE 2: AN EXAMPLE OF TAGS ASSOCIATED WITH A LECTURE

Knowledge removes inappropriate tags, as in accordance with our Terms of Service section 9, *Content Standards*.⁹ Furthermore, we are actively designing and developing a rating system that incentivizes the creation of high-quality, high-accuracy tags.

Tags are designed to be knowledge primitives — abstracted and decoupled from their parent lecture — they are use-agnostic and become the building blocks for endless educational products.

Playlists

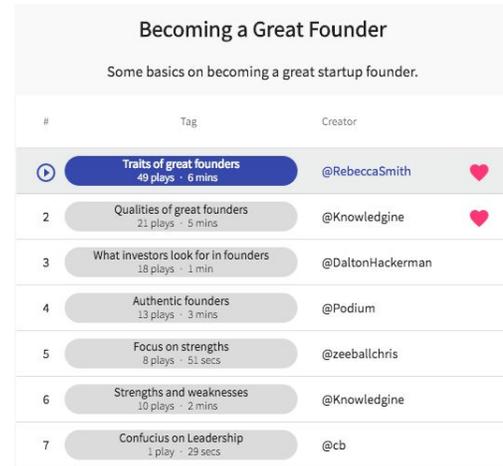
Playlists represent the first level of educational products that can be constructed with tags; a playlist is a collection of tags organized around a specific theme. Seen as an artistic form of creative curation for users, playlists effectively become micro-courses built on academic content. This novel system complements traditional learning techniques, while encouraging new content and professor discovery. Playlists can range from pragmatic applications in entrepreneurship (see figure 3) to exploratory applications such as “Top Psychology Professor Tags of State U”. The possibilities are endless when it comes to creating playlists; we envision a range of parties, from students, to small businesses, to working professionals, all utilizing and benefiting from them.

Playlists are designed to be searchable and shareable, allowing users on the platform to find and follow expertly crafted micro-courses on the subjects that matter to them most.

Profiles

Profiles are the next level of educational products that can be assembled through the use of tags. At a high level, profiles can be collections of tags, playlists, lectures, and courses all organized by the owner of the profile. A student who is interested in creating a biotech startup can curate playlists on both biotech and entrepreneurship. Other students who may also be interested in such a combination can follow that student (similar to other social media platforms) to continue learning.

Profiles can also represent non-profit organizations, societies, and institutions, encouraging content curation and dissemination related to key issues in society, politics, science, etc.



#	Tag	Creator
1	Traits of great founders 49 plays · 6 mins	@RebeccaSmith
2	Qualities of great founders 21 plays · 5 mins	@Knowledgegine
3	What investors look for in founders 18 plays · 1 min	@DaltonHackerman
4	Authentic founders 13 plays · 3 mins	@Podium
5	Focus on strengths 8 plays · 51 secs	@zeeballchris
6	Strengths and weaknesses 10 plays · 2 mins	@Knowledgegine
7	Confucius on Leadership 1 play · 29 secs	@cb

FIGURE 3: EXAMPLE OF A PLAYLIST RELATED TO STARTUPS

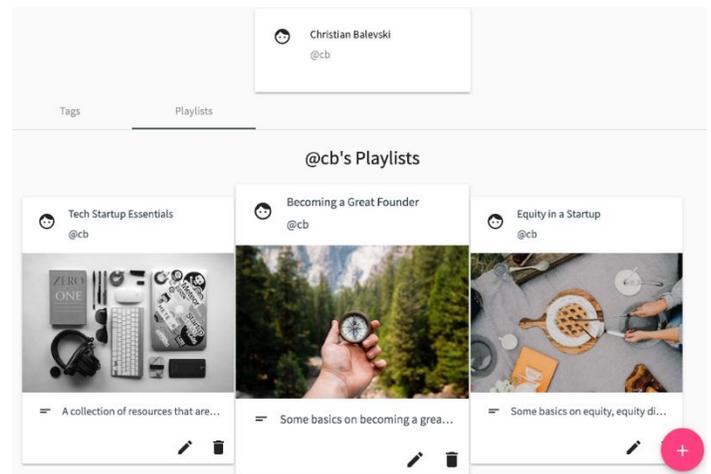


FIGURE 4: AN EXAMPLE OF A STARTUP FOCUSED PROFILE

Systems in Development or Testing

Beyond the organization systems that are currently implemented, Knowledgine has begun development of additional systems to enhance and secure the learning experience for users on our platform.

- **Entity Recognition:** An automated application of NLP used to identify relevant people, organizations, geopolitical entities, products, locations, and quantities to create more jumping-off points for students to learn and explore. We cross-reference Wikipedia data to disambiguate junk entities from legitimate ones.
- **Student Name Filter:** Student privacy is of utmost importance. If a professor mentions the first and last name of a student, our student name filter can identify it and remove it from the lecture.
- **Sentiment Analysis:** Understanding the sentiment of professor speech, cross-referenced with learning statistics, will give professors insights on how they can improve their lecture delivery.
- **Machine translation and speech synthesis:** Experimental application of speech synthesis where a professor's lecture is transcribed, translated (i.e. Spanish, French, Mandarin, etc.), and then spoken back in the respective language with the proper accent and attenuation using speech synthesis. This means a professor could give their lecture once, in one language, and learners everywhere could listen to it in their native tongue.

Content Distribution and User Experience

Knowledgine's content distribution strategy revolves around delivering the right piece of information to the right audience, regardless of their physical location.

- **Web Application:** Web-based interactive media player that allows students and users to engage with and access content. Listen to lectures, search and discover content, create tags, and curate playlists.
- **Search:** Search across tags, lectures, courses, and universities to allow learners to find exactly what they want.
- **Cloud Hosting:** Powered by Amazon Web Services (AWS) and MongoDB to offer redundant, secure, and globally available access to content.
- **Product Roadmap:** Recommendation Engine, Smart Speaker Integration (Amazon Alexa, Google Home, etc.), mobile, semantic search.

Web Application

Knowledgine's web application is the core location for user interaction on our platform. We offer a rich set of features that allows users to discover, curate, and experience academic knowledge in a modern and intuitive manner.

The web application is comprised of four main components (see figure 5 below):

- **Navigation Menu (left):** Quick access to controls that allow users to navigate around content to find what they want.
- **Main View (center):** The main display for information about courses, lectures, professors, universities, playlists, and search results.
- **Player Bar (bottom):** Houses the controls for lecture audio playback, speed, volume, queue, as well as accessing transcripts.
- **Tag Drawer (right):** Menu that houses the list of tags associated with a lecture, including the form to create new tags.

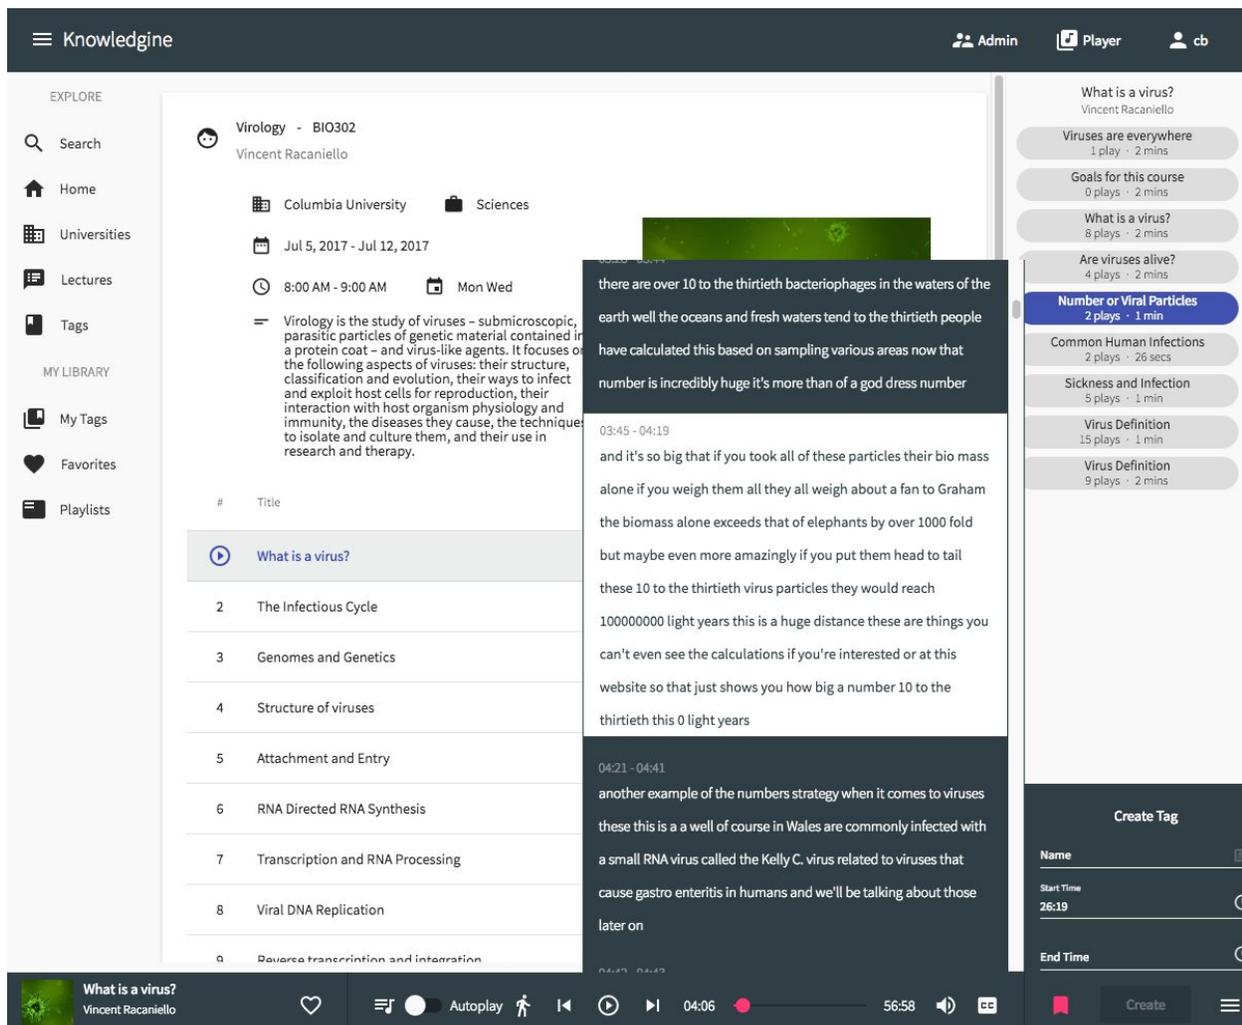
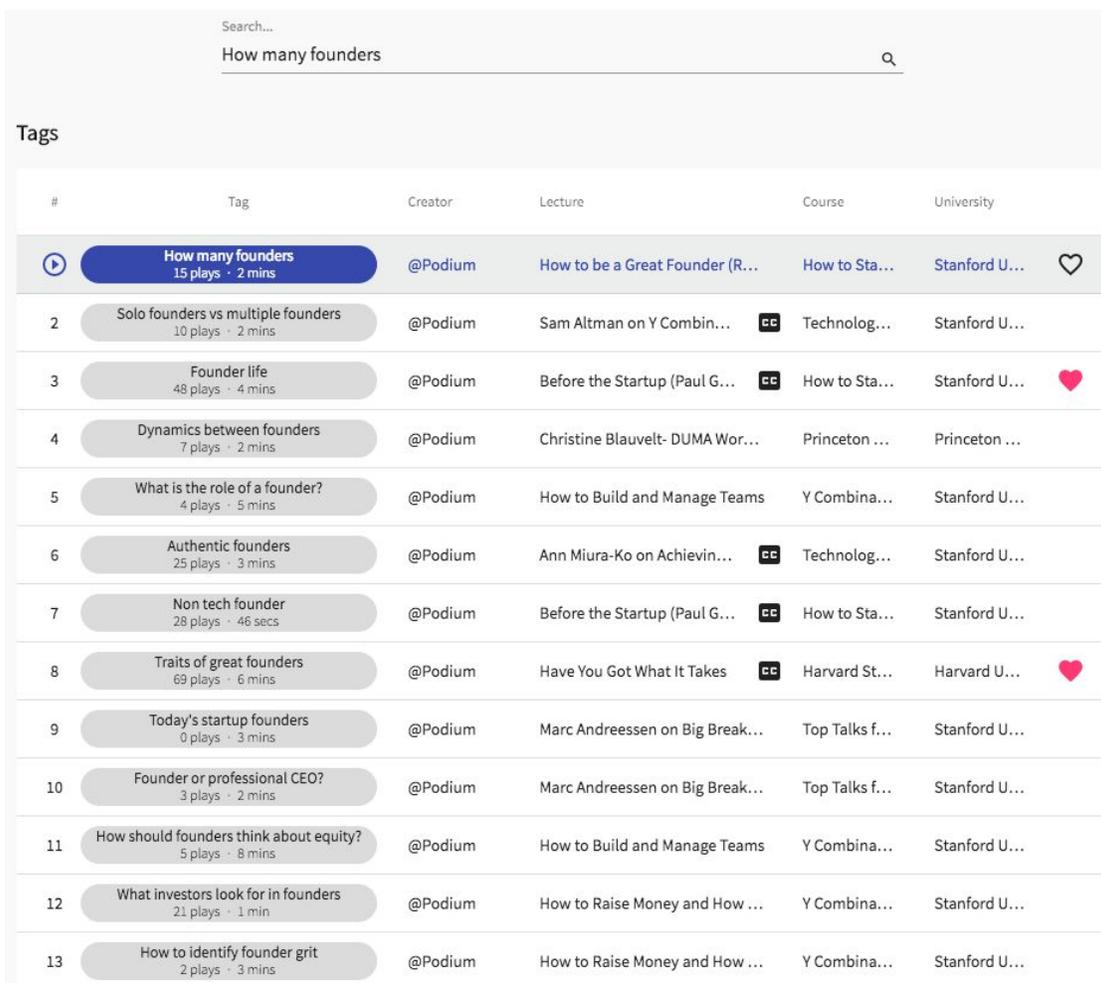


FIGURE 5: KNOWLEDGINE WEB APPLICATION

Search

Search functionality on Knowledgine is designed to allow users to find what they are interested in as fast as possible. To fulfill this objective, a query searches across every lecture from every university on our platform and returns a list of tags, lectures, and courses that are most closely related to the user's search parameters. The current search algorithm uses a MongoDB text search¹⁰ where indexes are created on various properties of tags, lectures, and courses. This system offers a working first iteration of a continuously improving search function. Currently, words are matched on their word stems rather than semantic meaning. An example of the type of results that this system produces can be found in figure 6 below. The important thing to note is that each tag result is a short clip explaining the exact subject that the user searched for. Users can quickly jump from one professor's explanation to another, developing a well-rounded and diverse understanding of the subject. Knowledgine believes this is a powerful tool for learning, and many of our users have reacted positively to this functionality.



The screenshot shows a search interface with a search bar containing the text "How many founders". Below the search bar, the word "Tags" is displayed. A table lists 13 search results, each with a rank number, a tag title, play count, duration, creator, lecture title, course, university, and a heart icon.

#	Tag	Creator	Lecture	Course	University
1	How many founders 15 plays · 2 mins	@Podium	How to be a Great Founder (R...	How to Sta...	Stanford U...
2	Solo founders vs multiple founders 10 plays · 2 mins	@Podium	Sam Altman on Y Combin...	Technolog...	Stanford U...
3	Founder life 48 plays · 4 mins	@Podium	Before the Startup (Paul G...	How to Sta...	Stanford U...
4	Dynamics between founders 7 plays · 2 mins	@Podium	Christine Blauvelt- DUMA Wor...	Princeton ...	Princeton ...
5	What is the role of a founder? 4 plays · 5 mins	@Podium	How to Build and Manage Teams	Y Combina...	Stanford U...
6	Authentic founders 25 plays · 3 mins	@Podium	Ann Miura-Ko on Achievin...	Technolog...	Stanford U...
7	Non tech founder 28 plays · 46 secs	@Podium	Before the Startup (Paul G...	How to Sta...	Stanford U...
8	Traits of great founders 69 plays · 6 mins	@Podium	Have You Got What It Takes	Harvard St...	Harvard U...
9	Today's startup founders 0 plays · 3 mins	@Podium	Marc Andreessen on Big Break...	Top Talks f...	Stanford U...
10	Founder or professional CEO? 3 plays · 2 mins	@Podium	Marc Andreessen on Big Break...	Top Talks f...	Stanford U...
11	How should founders think about equity? 5 plays · 8 mins	@Podium	How to Build and Manage Teams	Y Combina...	Stanford U...
12	What investors look for in founders 21 plays · 1 min	@Podium	How to Raise Money and How ...	Y Combina...	Stanford U...
13	How to identify founder grit 2 plays · 3 mins	@Podium	How to Raise Money and How ...	Y Combina...	Stanford U...

FIGURE 6: TAG SEARCH RESULTS

Cloud Hosting

Knowledgine is powered by Amazon Web Services (AWS) and MongoDB Cloud Atlas. AWS is the leader in cloud infrastructure and we use their services for everything from website hosting and data storage to content delivery. Knowledgine uses MongoDB as it's cloud database provider to offer a redundant and high-availability document storage system.

Product Roadmap

- **Recommendation Engine:** Content recommendation is an integral part of a media platform, and Knowledgine has been in communication with research scientists at NJIT to develop a “cold start” recommendation algorithm. Knowledgine has also begun researching a theoretical model for recommendation that could offer fine-grain search results based on individual learners.
- **Smart Speaker Integration:** “Alexa, play my last lecture...” Smart speakers are becoming more and more popular; integration with these devices will increase the ease of consumption of academic media.
- **Semantic Search:** Improve search algorithm to return results based on semantic relatedness and user intent rather than just word stem match.
- **API Integration:** Develop API integration to allow users to incorporate Knowledgine content in university/personal websites, LMS platforms, social media, etc.
- **Mobile:** Knowledgine is working on mobile compatibility and a responsive experience to allow users to learn on the go.

Professor Intellectual Property

Professor intellectual property (IP) is a core consideration of Knowledgine. We believe that professors should own their IP and be able to profit off of their work in a similar fashion to publications. Often times, professors are not fully aware of their university's rules and policies of IP ownership when it comes to course material. The exact policies of each university vary from school to school but the American Association of University Professors (AAUP) outlines a concise argument for IP ownership under section III paragraph B¹¹:

The "work-for-hire" doctrine is a statutory exception to the general ownership provisions of the copyright law. It is a way of allocating whether an employee or an employer is the author, and thus copyright holder, of work performed in the course of employment. The work-for-hire provision entitles an employer to assert ownership over materials prepared by its employees acting within the "scope of their employment".

Generally, faculty scholarly work is not considered work-for-hire. "[I]t has been the prevailing academic practice to treat the faculty member as the copyright owner of works that are created independently and at the faculty member's own initiative for traditional academic purposes." *Statement on Copyright*, AAUP Policy Documents & Reports 182 (9th ed. 2001).

Administration ownership of faculty scholarly works, lecture notes and teaching materials would profoundly contradict the practices of the academic community. Work for hire doesn't fit, legally or policy-wise, into the academic scholarship arena.

Academic freedom requires that faculty be free to produce work reflecting their own views and theories — not those of administration or trustees. If all work belonged to the university, then its content would also have to be controlled or at least accepted by the university, which would vitiate any freedom of thought or inquiry.

Business Model

Knowledgine is free for professors and universities to use. Knowledgine will run advertisements to generate revenue from users, while offering paid subscriptions for enhanced features and usability. This business model is proven within a variety of different companies and markets; from music streaming (Spotify) and education technology (Duolingo), to media creation and distribution (YouTube, Netflix, etc.), a freemium based subscription model offers a sustainable, and validated approach to creating a profitable business.

Professors and universities will be a part of a revenue share to further encourage Knowledgine's use and promote excellent content. Students will have free, ad-free access to their university's content.

Conclusion

Up until this point, openly sharing university knowledge has been done well by only the most privileged universities. Academic knowledge is not just important for students but for the world. As Trivella and Dimitrios note in their research, “universities are the main formal means of creating, disseminating, and transferring knowledge, which is the key-factor for the growth of the global economy.”¹² Knowledgine offers universities and professors the ability to make this vision a reality. By reducing costs, automating lecture capture and organization, and distributing content to learners in a modern way, Knowledgine is rewriting the status quo of academic media. We're challenging systemic attachment to video, while becoming a bastion of open and accessible knowledge for everyone.

After all,

“Science knows no country, because knowledge belongs to humanity, and is the torch which illuminates the world.” - Louis Pasteur

References

1. MIT OCW - <http://web.mit.edu/fnl/volume/231/miyagawa.html>
2. MIT OCW Cost - <https://ocw.mit.edu/give/why-give/>
3. Cost of MOOCs - <http://www.irrodl.org/index.php/irrodl/article/view/1901/3069>
4. MOOC Statistics - <https://www.insidehighered.com/news/2013/03/08/researchers-explore-who-taking-moocs-and-why-so-many-drop-out>
5. Music Consumer Data - <https://www.nielsen.com/us/en/insights/news/2017/time-with-tunes-how-technology-is-driving-music-consumption.html>
6. Podcast Consumer Data - <http://www.edisonresearch.com/wp-content/uploads/2017/04/Podcast-Consumer-2017.pdf>
7. Saint Louis University of Missouri - <https://www.slu.edu/news/2018/august/slu-alexa-project.php>
8. IBM Speech Transcription - <https://www.ibm.com/blogs/watson/2017/03/reaching-new-records-in-speech-recognition/>
9. Knowledge's Terms of Use - <https://www.knowledgine.com/terms>
10. Search Algorithm - <https://docs.mongodb.com/manual/text-search/>
11. Academic IP Ownership - <https://www.aaup.org/issues/copyright-distance-education-intellectual-property/resources-copyright-distance-education-and/intellectual-property-issues-faculty>
12. Knowledge Management Strategy within Higher Education <https://www.sciencedirect.com/science/article/pii/S1877042815012872>