Latent Lab: Large Language Models for Knowledge Exploration

Trudy Painter[1], Kevin Dunnell[2], Andrew Stoddard[1], Andy Lippman[2]
[1]Department of Electrical Engineering and Computer Science, MIT
[2]MIT Media Lab
{tpainter, dunnell, apstodd, lip}@mit.edu

Abstract

This paper investigates the potential of AI models, particularly large language models (LLMs), to support knowledge exploration and augment human creativity during ideation. We present "Latent Lab" an interactive tool for discovering connections among MIT Media Lab research projects, emphasizing "exploration" over search. The work offers insights into collaborative AI systems by addressing the challenges of organizing, searching, and synthesizing content. In a user study, the tool's success was evaluated based on its ability to introduce users to an unfamiliar knowledge base, ultimately setting the groundwork for the ongoing advancement of human-AI knowledge exploration systems.

Introduction

The untapped potential of collective knowledge holds significant implications for idea evolution and innovation across various entities (Curley and Salmelin 2013). Despite the digital revolution, information organization remains strikingly similar to traditional methods, limiting exploration across diverse sources and impeding the discovery of interconnected relationships. Current search approaches prioritize quick answers and display results in a list format. This hinders the discovery of interconnected relationships required for meaningful exploration and undermines the context of search terms by prioritizing keywords over semantics.

In contrast, synthesis tools like ChatGPT[1] offer a paradigm shift in user interface design through conversational interaction, though they have drawbacks such as the opaqueness of information sources and limited text-based interaction. This paper outlines the development of Latent Lab[2] and evaluates it in the context of the MIT Media Lab data set of 4,000+ research projects. This exploration tool surpasses previous search and synthesis tools by incorporating scanning and active visual interaction. Leveraging data manipulation libraries, interactive visuals, and LLMs, Latent Lab overcomes the constraints of keyword-centric search, allowing users to engage in semantically meaningful exploration and synthesis across diverse data sets. The iterative design process of the tool itself highlights the importance of exploration in the creative process, offering a glimpse into the potential of AI-assisted idea generation.

We make the following contributions to the field of human-AI interaction systems.

- We present the design and implementation of an interactive knowledge visualization tool, including a novel automated technique to label idea clusters using an LLM.
- We report the results from a user evaluation study, demonstrating the utility of a hybrid search/synthesis system to find meaningful insights and connections often missed by traditional search and synthesis tools.

Related Work

Knowledge Organization

Vannevar Bush's memex laid the foundation for hypertext and associative indexing (Bush 1945). Richard Feynman's triangulation method emphasized understanding relationships between concepts (Feynman, Gottlieb, and Leighton 2006). These ideas influenced the development of Google Knowledge Graph (Carr 2007). Our approach to knowledge organization builds on these works to enable fluid exploration of linked information.

Information Visualization

Shneiderman's taxonomy established information visualization principles, with the "overview, zoom and filter, details-on-demand" mantra guiding the design of visual interfaces for interacting with large data sets. (Shneiderman 1996). Bostock et al. presented D3.js for interactive visualizations (Bostock, Ogievetsky, and Heer 2011). Heer and Shneiderman highlighted the importance of interaction in visual analysis (Heer and Shneiderman 2012). Our work integrates these principles to create an informative interface for users.

Information Retrieval

Spärck Jones introduced the tf-idf weighting scheme for keyword-based search (Spärck Jones 1972). Mikolov et al. proposed the Word2Vec model for embedding-based search (Mikolov et al. 2013). Devlin et al. developed the BERT model, which further improved semantic search (Devlin et al. 2018). Latent Lab extends this work by leveraging embedding-based search to provide users with relevant results in complex information landscapes.

[1]`https://chat.openai.com/`

[2]Try Latent Lab at `https://latentlab.ai/`

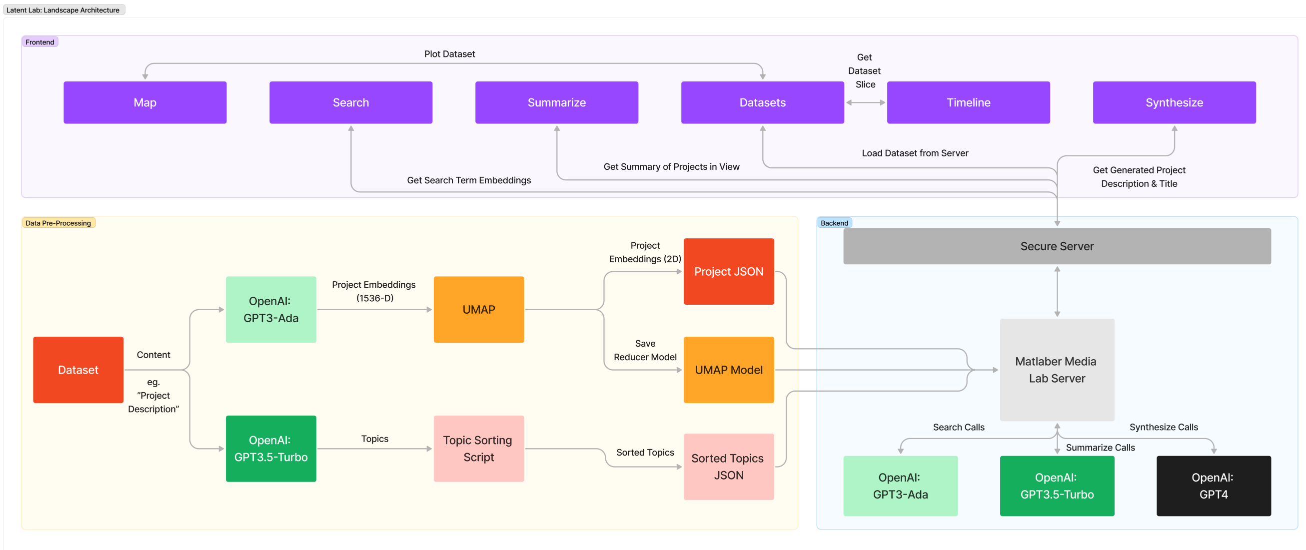

Figure 1: System Architecture of Latent Lab

Human-AI Collaboration

Minsky's Society of Mind proposed human intelligence as a result of interacting agents (Minsky 1988). Influential works that consider humans and intelligent systems as interacting agents include TRIZ, Polya's work on invention, and Weis and Jacobson's DELPHI framework (Weis and Jacobson 2021; Polya 1945; Altshuller 1999). Our work further examines human-AI collaboration, aiming to create a system that amplifies human capabilities and positions AI as a "copilot" rather than an "autopilot."

Methods

System Overview

Latent Lab is an AI-powered knowledge exploration system. High-dimensional unstructured data is condensed and visualized in an interactive 2D map. The interface allows users to explore labeled clusters of similar topics, search by semantic context, and synthesize new ideas.

System Architecture

The system, designed for performance and scalability, utilizes Fast API[3] and Python, for the back end, while the front end employs Vercel,[4], Next.js[5], React[6], and TypeScript. The initial approach was to execute all operations on the front end. However, due to the absence of a fully JavaScript-ported version of UMAP (McInnes, Healy, and Melville 2018), we decided to incorporate a back end server. This also facilitated server-side rendering to expedite data set loading. See Figure 1 for the system architecture diagram.

[3]https://fastapi.tiangolo.com/
[4]https://vercel.com
[5]https://nextjs.org
[6]https://reactjs.org

Data Processing

The data processing pipeline is mostly automated and runs independently of the web app back end for each new data set. It generates three primary artifacts:

- A project JSON containing the unstructured data and embedding data for mapping every project on the front end
- A sorted research topics JSON containing all topics produced by the pipeline, ordered by topics with the most associated projects, used for the labels on the front end
- A pickled UMAP model to reduce project and topic embeddings to 2 dimensions on the back end

Topic Extraction

Latent Lab's automated topic extraction feature sets it apart from other embedding visualization tools, which don't provide insights into cluster meanings. The system uses GPT-3.5-Turbo to distill topics for each project, count occurrences of unique topic labels, and identify related projects. It then calculates label positions using the centroid of the UMAP-reduced coordinates for each associated topic.

Components

The Latent Lab interface has four main components, shown in Figure 2. It includes a Map Visualization, Generation Workbench, Search Bar, and Timeline Slider.

Map Visualization

The main visualization displays an organized map of project data, with dots representing research projects and clusters indicating semantic similarity. Dot colors correspond to different Media Lab groups and can be customized to represent other discrete data set attributes.

Contour lines in the map indicate data density within clusters, a concept borrowed from topographic maps where they

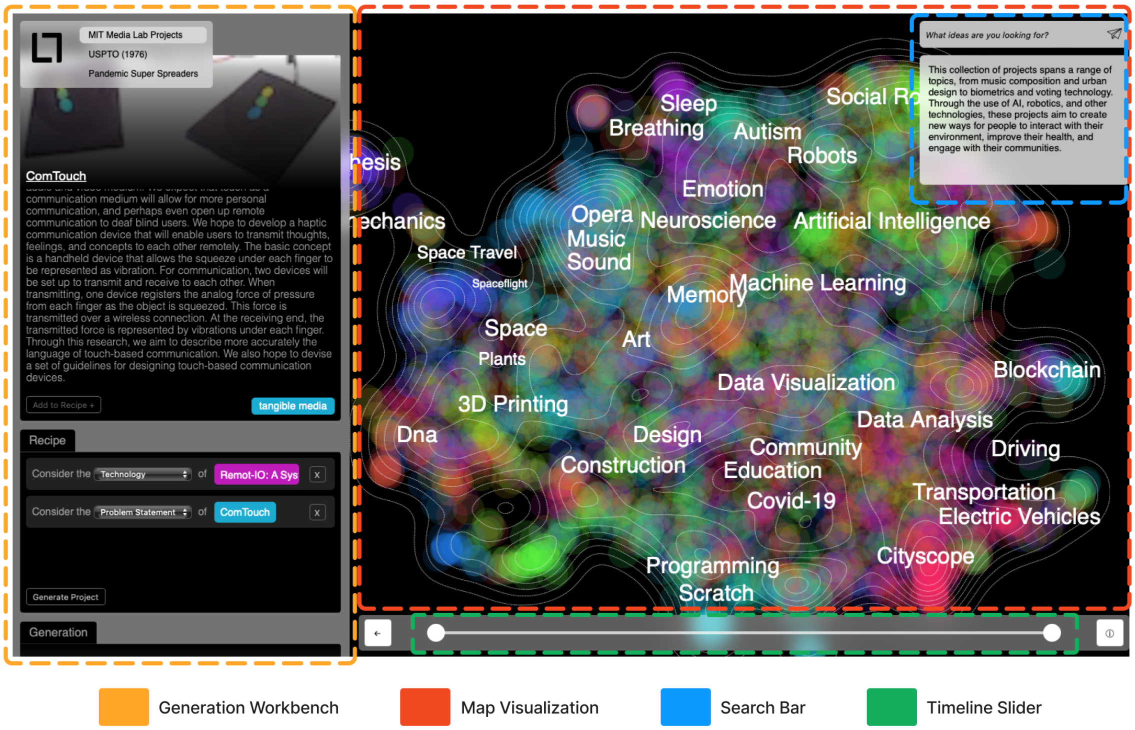

Figure 2: Latent Lab Interface, Annotated to Differentiate Between Components

represent elevation. Paired with the timeline, the changing contour lines reveal the evolution of research concentration.

Users can pan and zoom, uncovering varying levels of information. High-level labels and contour lines are shown at the highest zoom level, while sub-topic labels and project details appear when zooming in. An occlusion algorithm determines label visibility based on popularity and bounding box overlap.

Generation Workbench

Latent Lab's Generation Workbench allows users to create a "recipe" for collaboratively synthesizing new research project ideas. Users can choose whole projects or specific aspects, such as community, problem statement, or technology, to include. After designing the recipe, selecting "generate" submits a preset prompt with selected project elements to GPT-4 via the OpenAI API, producing a synthesized project title and description. Users can view the exact prompt by clicking the "What was used to generate this?" information button. See Figure 3 for the user flow diagram.

Search & Summarization

Latent Lab employs embedding-based search for semantic meaning instead of simple keyword-matching, enabling more intuitive project exploration through contextual relationships. When a user searches, the query is sent to the back end server, and the GPT-Ada API returns a 1,536-value embedding. This is passed to the UMAP reducer, yielding x and y coordinates, which are sent to the front end to dynamically zoom and highlight the relevant map region. Figure

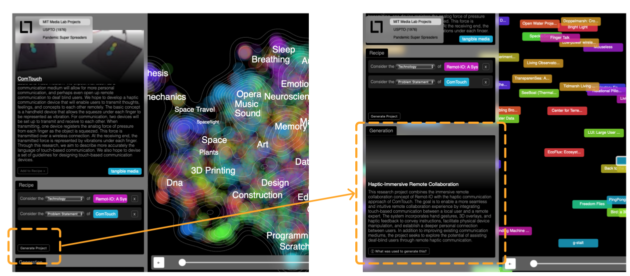

Figure 3: Generating a Research Project Idea

4 demonstrates this process using "quadratic voting" as the search term. Below the search bar, Latent Lab displays summaries that provide users with a quick overview of projects in the currently viewed map region.

Timeline Slider

Latent Lab's Timeline feature enables users to explore data set progression over a selected period using start and end date sliders. This functionality, particularly useful alongside the search bar, allows for efficient examination of current or ongoing projects in specific areas. Figure 5 illustrates timeline filtering for projects since 2018.

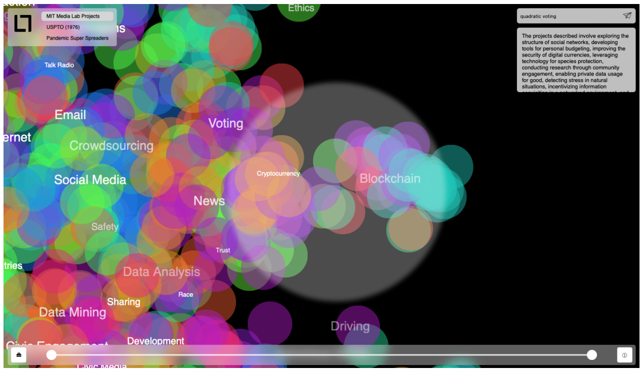

Figure 4: Search Highlighting in Map

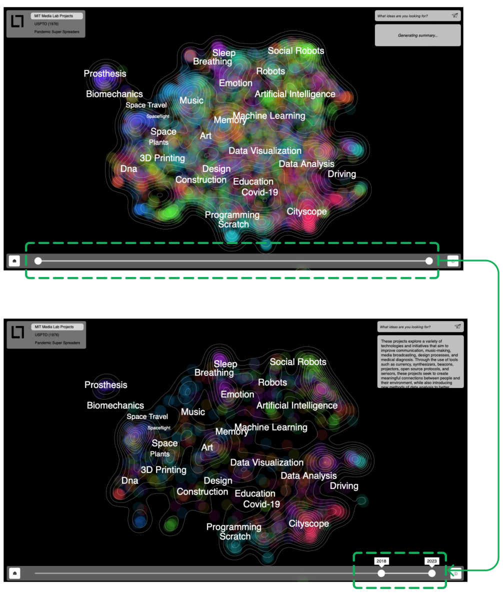

Figure 5: Timeline Evolution

Study Overview

We designed a study to evaluate users' experience exploring MIT Media Lab research using Latent Lab. The study compared Latent Lab to the current MIT Media Lab website, which uses traditional keyword-based search. Surveying 94 researchers via Prolific, participants interacted with both tools in a randomized order. After using each tool, they answered questions assessing clarity, effort (Hart and Staveland 1988), engagement, mental support, future use, trust (benevolence, capability, and reliability) (Mayer, Davis, and Schoorman 1995), and insight on a 1-5 Likert scale. The study aimed to measure Latent Lab's effectiveness in fostering human-AI collaboration, enhancing user experience, and promoting a deeper understanding of Media Lab projects when users engage with AI-powered tools.

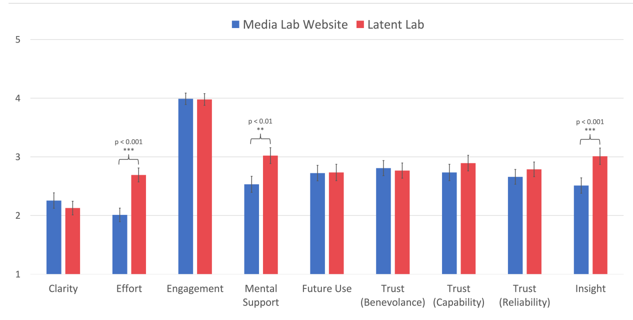

Figure 6: User Evaluation Results

Results

Analysis

The study results in Figure 6 indicate that Latent Lab shows promise as an AI-assisted exploration tool, compared to the Media Lab website. Participants were equally engaged, trusted both systems, and expressed equal likelihood to use them in the future.

Although Latent Lab required more effort, this is likely due to its novel functionality compared to traditional search interfaces. As users become more familiar with the design, we expect the effort to decrease, facilitating seamless human-AI collaboration. Latent Lab outperformed the Media Lab website in providing higher mental support and insight, suggesting that its semantic map effectively organizes knowledge and offers a deeper understanding of MIT Media Lab research than the current Media Lab website.

Overall, the study highlights Latent Lab's potential and underscores the need for minor improvements to deliver a consistent user experience and enhanced search results.

Initially motivated by the ambition to generate research project ideas, our work was reoriented by early user feedback towards enhancing knowledge organization for optimized exploration. Time constraints precluded us from thoroughly evaluating generated ideas' utility, novelty, and feasibility. Our ultimate objective is to probe and benchmark the creative boundaries of Latent Lab as a human-AI co-invention system.

Latent Lab has already attracted interest from research groups at MIT and Media Lab sponsor companies. In response, our focus has shifted again, this time towards optimizing the system to support tens of millions of data points and simplifying the data upload process for users, enabling them to seamlessly integrate their organization's data for deeper insights.

Conclusion

Latent Lab serves as an innovative and powerful tool for exploring interconnected relationships within diverse data sets. By utilizing LLMs and visually engaging interfaces, it transcends conventional search limitations, providing a semantically meaningful and context-aware experience. Emphasizing the value of exploration and iterative design, Latent

Lab realizes the long-sought goal of information technology experts for an intuitively accessible wealth of interconnected information. AI-assisted exploration has turned this vision into reality, setting the stage for future human-AI co-invention systems and fostering more intuitive and productive collaborations that are capable of generating novel and impactful creations.

Author Contributions

TP and KD are leading authors. TP, KD, and AL contributed to the conception of Latent Lab. TP, KD, and AP all contributed to the development of Latent Lab.

Acknowledgments

We thank Joshua Bello, Deborah Jang, Michael Peng, Jinhee Won, and Anushka Nair for their help with Latent Lab development. We also thank Ziv Epstein and Hope Schroeder for their guidance with user testing. Finally, we thank all survey participants for their time.

References

[Altshuller 1999] Altshuller, G. 1999. *The Innovation Algorithm: TRIZ, systematic innovation, and technical creativity*. Technical Innovation Center, Inc.

[Bostock, Ogievetsky, and Heer 2011] Bostock, M.; Ogievetsky, V.; and Heer, J. 2011. D3: Data-driven documents. *IEEE transactions on visualization and computer graphics* 17(12):2301–2309.

[Bush 1945] Bush, V. 1945. As we may think. *The Atlantic Monthly* 176(1):101–108.

[Carr 2007] Carr, N. 2007. Freebase: A case study.

[Curley and Salmelin 2013] Curley, M., and Salmelin, B. 2013. Open innovation 2.0: a new paradigm. *OISPG White Paper* 1–12.

[Devlin et al. 2018] Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

[Feynman, Gottlieb, and Leighton 2006] Feynman, R. P.; Gottlieb, M. A.; and Leighton, R. 2006. *Feynman's Tips on Physics: A Problem-Solving Supplement to the Feynman Lectures on Physics*. New York: Basic Books.

[Hart and Staveland 1988] Hart, S. G., and Staveland, L. E. 1988. Development of nasa-tlx (task load index): Results of empirical and theoretical research. *Advances in psychology* 52:139–183.

[Heer and Shneiderman 2012] Heer, J., and Shneiderman, B. 2012. Interactive dynamics for visual analysis. *Communications of the ACM* 55(4):45–54.

[Mayer, Davis, and Schoorman 1995] Mayer, R. C.; Davis, J. H.; and Schoorman, F. D. 1995. An integrative model of organizational trust. *The Academy of Management Review* 20(3):709–734.

[McInnes, Healy, and Melville 2018] McInnes, L.; Healy, J.; and Melville, J. 2018. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.

[Mikolov et al. 2013] Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*.

[Minsky 1988] Minsky, M. 1988. *The Society of Mind*. Simon and Schuster.

[Polya 1945] Polya, G. 1945. *How to Solve It: A New Aspect of Mathematical Method*. Princeton University Press.

[Shneiderman 1996] Shneiderman, B. 1996. The eyes have it: A task by data type taxonomy for information visualizations. 336–343.

[Spärck Jones 1972] Spärck Jones, K. 1972. A statistical interpretation of term specificity and its application in retrieval. *Journal of Documentation* 28(1):11–21.

[Weis and Jacobson 2021] Weis, D., and Jacobson, R. 2021. Delphi: A framework for human-ai collaborative problem solving and learning. *arXiv preprint arXiv:2106.03852*.