

ResearchCube: Multi-Dimensional Trade-off Exploration for Research Ideation

ANONYMOUS AUTHOR(S)

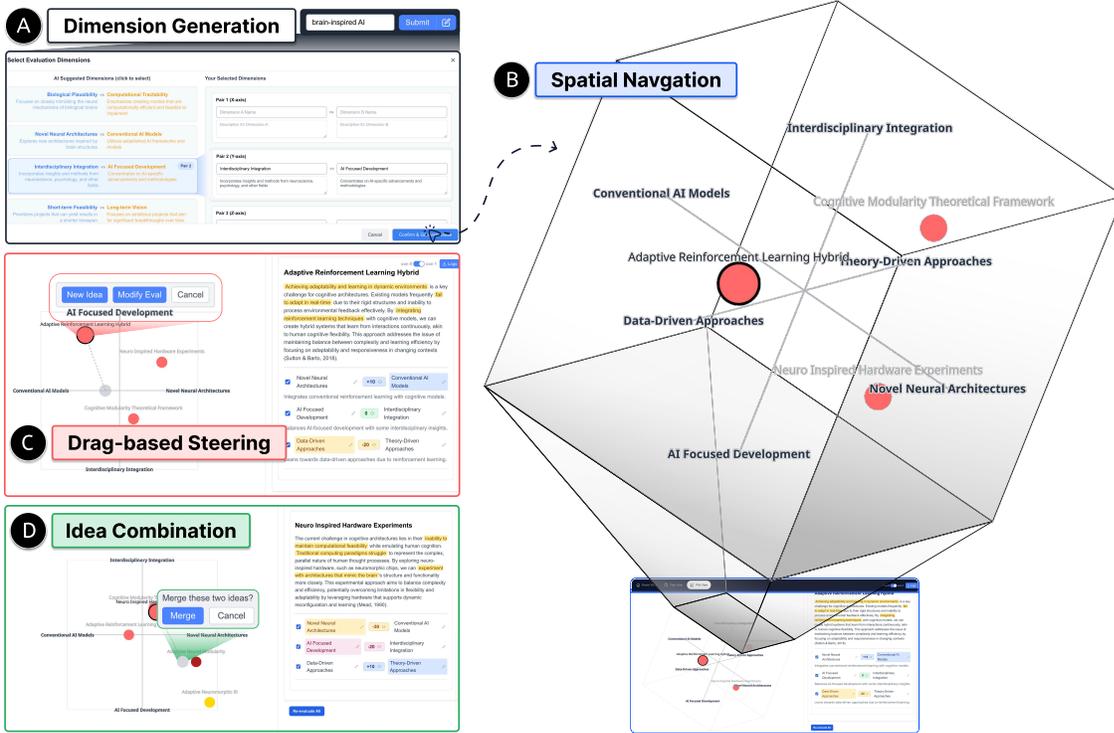


Fig. 1. ResearchCube renders research ideas as interactive nodes in a 3D evaluation space. Each axis represents a user-selected bipolar trade-off dimension (e.g., theory-driven vs. data-driven). Users explore ideas by rotating the cube, and refine them by dragging nodes toward desired trade-off balances.

Research ideation requires navigating trade-offs across multiple evaluative dimensions—yet most AI-assisted research ideation tools present a missed opportunity: none explicitly design for shared human-AI exploration of multi-dimensional trade-off spaces. This paper presents ResearchCube, a multi-dimensional trade-off exploration system that renders research ideas as manipulable nodes along user-selected bipolar dimensions (e.g., theory-driven vs. data-driven). Rather than unipolar scales where “more is better,” ResearchCube frames each axis as a trade-off spectrum with meaningful poles at both ends, making tensions in research choices explicit. To address the cold-start problem, the system proposes candidate dimension pairs from which users select up to three to construct a personalized 3D evaluation space. Four primary interactions—dimension generation, spatial navigation, drag-based steering, and idea combination—enable researchers to explore, compare, and refine ideas through direct spatial manipulation rather than textual prompts.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2026 Copyright held by the owner/author(s).

Manuscript submitted to ACM

53 Additional Key Words and Phrases: Research Ideation, Direct Manipulation, Multi-Dimensional Visualization, Generative AI, Intent-
54 based Interfaces

55
56 **ACM Reference Format:**

57 Anonymous Author(s). 2026. ResearchCube: Multi-Dimensional Trade-off Exploration for Research Ideation. In *CHI Conference*
58 *on Human Factors in Computing Systems (CHI '26)*, April 13–17, 2026, Barcelona, Spain. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.xxxx/xxxxxxx.xxxxxxx>
59
60

61
62
63 **1 INTRODUCTION**

64
65 Research ideation involves navigating a vast space of possibilities along multiple evaluative dimensions—novelty,
66 feasibility, impact, methodological rigor, and domain fit, among others. Navigating this multi-dimensional space is part
67 of what makes for creative ideas, particularly for “wicked,” open-ended problems with no obvious global optimum but
68 rather negotiations and trade-offs. While we have good systems for exploring ideas based on facets [9, 10] or spatial
69 manipulation [1, 12], evaluative multi-dimensionality is seldom an explicit focus—dimensionality is often left implicit
70 in favor of foregrounding idea facets or linear chats. This lack of explicit support leaves important and effortful work
71 to be done primarily using the researcher’s internal cognitive resources, missing opportunities for deeper human-AI
72 collaboration on research ideation.
73
74

75 Direct Intent Manipulation (DIM) [3] demonstrated that drag-based steering in a 2D canvas helps researchers iterate
76 on ideas more fluidly than prompt-only interfaces. However, DIM and similar systems evaluate ideas on fixed unipolar
77 dimensions (e.g., novelty scored 0–100), implicitly treating each criterion as “more is better.” Research evaluation is
78 rarely so straightforward: high novelty often comes at the cost of feasibility; broad scope may sacrifice depth. A more
79 nuanced approach would reframe each axis as a trade-off spectrum with meaningful poles at both ends—for example,
80 theory-driven vs. data-driven, or narrow domain vs. broad generalization.
81

82 This paper addresses this gap by introducing **ResearchCube**¹, a system for multi-dimensional trade-off exploration
83 that renders research ideas as manipulable nodes along user-selected bipolar dimensions. This bipolar framing makes
84 trade-offs explicit and supports more nuanced positioning of ideas. To help users construct their evaluation space,
85 ResearchCube provides AI-assisted trade-off dimension generation: given a research intent, the system proposes
86 candidate bipolar dimension pairs, from which users select up to three to define the axes of an evaluation space (up to
87 3D). This scaffolding addresses the cold-start problem—users who lack clear evaluative frameworks can bootstrap from
88 AI suggestions while retaining full control over which trade-offs matter for their context.
89
90

91 This paper addresses the following research questions:

- 92
93
- 94 • **RQ1.** How can we represent research ideas in a multi-dimensional space where each axis corresponds to a
95 bipolar trade-off dimension rather than a unipolar scale?
 - 96 • **RQ2.** How can direct manipulation interactions (drag-to-steer, drag-to-merge) be mapped onto this trade-off
97 space to support idea iteration and synthesis?
- 98

99 This paper presents a working system that instantiates these design ideas, demonstrating how trade-off-based
100 spatialization combined with direct manipulation can support fluid exploration of research idea spaces.
101

102
103 ¹Open-source implementation available at: [https://github.com/\[anonymized\]](https://github.com/[anonymized]) (link will be provided upon acceptance).
104

Table 1. Feature comparison across related systems that support spatial representation and AI-assisted idea generation.

System	Dimension Generation	Trade-off Dimensions	Spatial Navigation	Drag-Based Steering	Idea Combination	Idea Fragmentation
Sensecape (UIST'23) [13]						
Graphologue (UIST'23) [8]						
Luminate (CHI'24) [12]	✓					
PatchView (UIST'24) [1]		✓		✓		
Toyteller (CHI'25) [2]				✓		
IdeaSynth (CHI'25) [9]					✓	
Scideator (CHI'25) [10]					✓	
DIM (UIST'25 Poster) [3]				✓	✓	
ResearchCube	✓	✓	✓	✓	✓	✓

2 RELATED WORK

Direct manipulation—continuous representation of objects, physical actions instead of syntax, and rapid reversible operations [11]—has proven effective for complex cognitive tasks [7]. Spatial reasoning provides an intuitive substrate for expressing nuanced distinctions [5], as humans naturally think in terms of proximity, direction, and relative position. Semantic interaction systems like ForceSPIRE [4] leverage this capacity by mapping drag gestures to updates in underlying analytical models, allowing users to steer document clustering through spatial arrangement. Recent work applies these principles to generative AI: PatchView [1] enables drag-based manipulation along opposing dimensions for worldbuilding, and Direct Intent Manipulation (DIM) [3] showed that dragging ideas toward target positions can steer research ideation more fluidly than prompt-only interfaces.

However, direct manipulation of ideas requires a dimensional space in which to manipulate them. This raises the question of how to construct meaningful evaluation dimensions. Luminate [12] addresses this through AI-assisted dimension generation, proposing attribute dimensions based on user-provided examples for design exploration. Yet Luminate’s dimensions emphasize attribute combinations rather than explicit trade-offs. Similarly, Sensecape [13] enables multi-level abstraction of LLM-generated content and Graphologue [8] visualizes outputs as interactive node-link diagrams, but neither provides evaluative dimensions that users can directly manipulate to steer generation.

For research ideation specifically, recent systems support iterative idea development through various mechanisms. IdeaSynth [9] enables facet-based idea combination, allowing researchers to merge concepts from different sources. Scideator [10] supports scientific ideation through structured recombination of research components. These systems help researchers explore and refine ideas, but they do not employ spatial representations for steering, nor do they make evaluative trade-offs explicit. The dimensional structure that could organize ideas spatially remains implicit, leaving researchers to navigate trade-offs through their internal cognitive resources.

ResearchCube bridges this gap by combining direct manipulation with AI-assisted dimension generation tailored for research ideation. Unlike prior systems that treat dimensions as independent attributes, ResearchCube frames each axis as a bipolar trade-off with meaningful poles at both ends, making tensions visible and manipulable. This work treats evaluative dimensions as both organizational scaffolds and control surfaces: the dimensional structure that organizes ideas also serves as the interface for modifying them. Table 1 summarizes how ResearchCube compares with related work across representative features.

3 SYSTEM DESIGN

ResearchCube enables researchers to explore, evaluate, and refine research ideas through direct spatial manipulation. The system renders ideas as interactive nodes in a multi-dimensional evaluation space, where each axis represents a bipolar trade-off dimension selected by the user.

3.1 Interaction Design

ResearchCube supports four primary interactions: (1) *dimension generation*, where users construct their evaluation framework from AI-suggested trade-off pairs; (2) *spatial navigation*, enabling 3D rotation with camera snapping for precise manipulation; (3) *drag-based steering*, allowing users to reposition ideas toward new trade-off balances; and (4) *idea combination*, synthesizing multiple ideas through proximity-based merging or fragment incorporation.

3.1.1 Dimension Generation. A central challenge in multi-dimensional ideation is that users often lack a clear evaluative framework before they begin exploring. While Luminare [12] introduced AI-assisted dimension generation for design exploration, ResearchCube extends this approach to research ideation with explicitly bipolar trade-off dimensions. To address this cold-start problem, the system generates candidate trade-off dimension pairs based on the user’s initial research intent. When a user enters an intent (e.g., “reinforcement learning for user intent prediction”), the system proposes five bipolar dimension pairs, each representing a meaningful trade-off spectrum relevant to the research domain. For example:

- Model Complexity: simple vs. complex
- Methodology: theory-driven vs. data-driven
- Scope: narrow domain vs. broad generalization
- Novelty: incremental vs. paradigm shift
- Validation: simulation vs. real-world deployment

The interface presents these suggestions in a two-column selection panel: the left column displays AI-generated dimension pairs with descriptions explaining each pole; the right column shows three axis slots (X, Y, Z) where users assign their chosen pairs. Users may select one, two, or three pairs depending on the desired complexity of exploration. Once dimensions are confirmed, the system generates diverse seed ideas that vary along the selected dimensions, ensuring broad coverage of the evaluation space. Each idea is scored and positioned according to its trade-off balance, with results streaming progressively into the space.

3.1.2 Spatial Navigation. Navigating a 3D space on a 2D screen presents interaction challenges: free rotation can disorient users, and dragging in arbitrary orientations creates ambiguity about which dimensions are being modified. The system addresses this through *camera snapping*. Users can freely rotate the evaluation cube to inspect ideas from any angle. However, when rotation stops, the camera automatically snaps to the nearest orthogonal face—front, back, left, right, top, or bottom—aligning the view with exactly two axes while locking the third (depth) axis. This snapping ensures that subsequent drag operations map unambiguously to the two visible dimensions. Users can also explicitly toggle individual dimensions on or off: disabling one dimension collapses the 3D cube to a 2D plane; disabling two produces a 1D line. This scalable dimensionality lets users focus on specific trade-offs or examine the full multi-dimensional landscape as needed.

209 3.1.3 *Drag-Based Steering*. The central interaction paradigm is *drag-based steering*: users express how they want an
210 idea to change by dragging its node to a new position in the evaluation space. When a user begins dragging, a ghost
211 node connected by a dashed line visualizes the proposed trajectory from the original position to the target. This preview
212 helps users understand exactly how the idea’s scores will change before committing. Upon releasing the drag, the
213 system offers two options: (1) **Generate New Idea**—the system rewrites the idea’s content to match the target position,
214 producing a variant that embodies the new balance of trade-offs; or (2) **Re-evaluate Only**—the system updates the
215 idea’s scores without changing content, useful when users disagree with the AI’s initial scoring. The steering interaction
216 makes intent expression spatial and continuous rather than verbal and discrete, supporting rapid iteration.
217
218
219

220 3.1.4 *Idea Combination*. Beyond refining individual ideas, users often want to synthesize promising elements from
221 multiple sources. The system supports two combination mechanisms at different granularities:
222

223 **Drag-to-merge**. When a user drags one idea node close to another (within a proximity threshold), the target node
224 highlights to indicate a potential merge. Releasing the drag triggers synthesis: the system combines the contributions of
225 both parent ideas into a new hybrid idea that integrates their complementary strengths. The merged idea is then scored
226 along all active dimensions and positioned accordingly in the evaluation space.
227

228 **Fragment incorporation**. Sometimes users identify a valuable phrase or concept within an idea that they want
229 to incorporate elsewhere without adopting the entire idea. Users can extract text snippets from any idea and save
230 them as fragments. To incorporate a fragment, users drag it onto any idea node; the system then revises the target idea
231 by integrating the fragment’s content. This supports non-linear ideation where users collect promising components
232 throughout a session and recombine them flexibly.
233
234
235

236 3.2 Design Implications

237 The design and development of ResearchCube surfaces several implications for future AI-assisted ideation systems.

238 **Spatialization as shared representation**. By rendering ideas as positions in a multi-dimensional evaluation
239 space, ResearchCube creates a *shared representation* [6] that both humans and AI can read and write. Users inspect
240 the landscape to understand how ideas compare; AI uses the same coordinate system to score new ideas and interpret
241 steering commands. This bidirectional legibility reduces the “black box” quality of LLM-based generation: users see
242 why an idea landed where it did (via per-dimension scores) and can contest that placement by dragging. Future systems
243 may extend this principle to other creative domains—design exploration, policy analysis, or scientific hypothesis
244 generation—wherever trade-offs can be externalized as interpretable axes.
245
246

247 **AI-scaffolded evaluation frameworks**. A persistent challenge in open-ended ideation is that users often lack a
248 clear evaluative framework before they begin. ResearchCube addresses this by having AI propose candidate dimension
249 pairs, lowering the barrier to entry while preserving user agency (users select which dimensions to adopt and can
250 edit labels). This scaffolding strategy—AI suggests structure, users curate and customize—may generalize to other
251 sensemaking tasks where the relevant criteria are not known in advance.
252

253 **Dimensionality as a design variable**. Supporting seamless transitions across dimensionality levels (up to 3D)
254 revealed that dimensionality itself is a useful design lever. Early-stage exploration benefits from lower dimensionality
255 (fewer axes, simpler visualization); as users develop richer mental models, they can “unlock” additional dimensions to
256 examine finer trade-offs. Systems that fix dimensionality upfront may miss opportunities to match interface complexity
257 to user readiness.
258
259
260

REFERENCES

- [1] John Joon Young Chung and Max Kreminski. 2024. Patchview: LLM-powered Worldbuilding with Generative Dust and Magnet Visualization. In *Proceedings of the 37th Annual ACM Symposium on User Interface Software and Technology (UIST '24)*. ACM. <https://doi.org/10.1145/3654777.3676352>
- [2] John Joon Young Chung, Melissa Roemmele, and Max Kreminski. 2025. Toyteller: AI-powered Visual Storytelling Through Toy-Playing with Character Symbols. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM. <https://doi.org/10.1145/3706598.3713435>
- [3] Zijian Ding, Fenghai Li, Haofei Yu, and Joel Chan. 2025. Towards Direct Intent Manipulation: Drag-Based Research Ideation, Evaluation and Evolution. In *The 38th Annual ACM Symposium on User Interface Software and Technology (UIST Adjunct '25)*. ACM, New York, NY, USA, 1–3. <https://doi.org/10.1145/3746058.3758453>
- [4] Alex Endert, Patrick Fiaux, and Chris North. 2012. Semantic Interaction for Visual Text Analytics. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI '12)*. ACM, New York, NY, USA, 473–482. <https://doi.org/10.1145/2207676.2207741>
- [5] Peter Gardenfors. 2004. *Conceptual spaces: The geometry of thought*. MIT press.
- [6] Jeffrey Heer. 2019. Agency plus automation: Designing artificial intelligence into interactive systems. *Proceedings of the National Academy of Sciences* 116, 6 (Feb. 2019), 1844–1850. <https://doi.org/10.1073/pnas.1807184115>
- [7] Edwin L. Hutchins, James D. Hollan, and Donald A. Norman. 1985. Direct Manipulation Interfaces. *Human-Computer Interaction* 1, 4 (Dec. 1985), 311–338. https://doi.org/10.1207/s15327051hci0104_2 Publisher: Taylor & Francis _eprint: https://doi.org/10.1207/s15327051hci0104_2.
- [8] Peiling Jiang, Jude Rayan, Steven P. Dow, and Haijun Xia. 2023. Graphologue: Exploring Large Language Model Responses with Interactive Diagrams. <https://doi.org/10.1145/3586183.3606737> arXiv:2305.11473 [cs].
- [9] Kevin Pu, K. J. Kevin Feng, Tovi Grossman, Tom Hope, Bhavana Dalvi Mishra, Matt Latzke, Jonathan Bragg, Joseph Chee Chang, and Pao Siangliulue. 2025. IdeaSynth: Iterative Research Idea Development Through Evolving and Composing Idea Facets with Literature-Grounded Feedback. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM.
- [10] Marissa Radensky, Simra Shahid, Raymond Fok, Pao Siangliulue, Tom Hope, and Daniel S. Weld. 2025. Scideator: Human-LLM Scientific Idea Generation Grounded in Research-Paper Facet Recombination. In *CHI Conference on Human Factors in Computing Systems (CHI '25)*. ACM.
- [11] Ben Shneiderman. 1982. The future of interactive systems and the emergence of direct manipulation. *Behaviour & Information Technology* 1, 3 (1982), 237–256.
- [12] Sangho Suh, Meng Chen, Bryan Min, Toby Jia-Jun Li, and Haijun Xia. 2024. Luminare: Structured Generation and Exploration of Design Space with Large Language Models for Human-AI Co-Creation. In *Proceedings of the CHI Conference on Human Factors in Computing Systems*. 1–26. <https://doi.org/10.1145/3613904.3642400> arXiv:2310.12953 [cs].
- [13] Sangho Suh, Bryan Min, Srishti Palani, and Haijun Xia. 2023. Sensecape: Enabling Multilevel Exploration and Sensemaking with Large Language Models. In *Proceedings of the 36th Annual ACM Symposium on User Interface Software and Technology*. ACM, San Francisco CA USA, 1–18. <https://doi.org/10.1145/3586183.3606756>