Topic Evolution and Technology Opportunity Emergence in Embodied Intelligence: A Dual-Source Analysis of Publications and Patents

Yan Zeng[a]

School of Economics and Management, Beijing Jiaotong University, Beijing, China

[a]23241004@bjtu.edu.cn

Abstract

Embodied intelligence has become a major research frontier in artificial intelligence, yet the evolution of this field has rarely been examined through the combined lens of scientific publications and technological patents. This paper introduces a Topic Evolution and Technology Opportunity (TETO) framework to investigate how research themes evolve and how technology opportunities emerge across the two knowledge systems. A Dynamic Latent Dirichlet Allocation (LDA) model with temporal dependence is employed to characterize topic evolution over time, after which publication and patent trajectories are examined comparatively. The analysis reveals a persistent divergence between academic research and technological innovation. Scientific publications place greater emphasis on cognitive mechanisms and Human-Machine interaction, whereas patents are concentrated more heavily on motion control and system integration. At the same time, several research themes exhibit convergent evolutionary patterns, reflecting increasing interaction between scientific exploration and technological development. Based on these cross-domain dynamics, technology opportunities are classified into three categories: frontier-exploration, industry-driven, and co-evolutionary. The proposed framework extends existing approaches to technology forecasting and provides a systematic basis for strengthening alignment between academic research and industrial innovation.

Keywords

Embodied intelligence, dynamic LDA, topic evolution, publication-patent dual-source analysis, technology opportunity identification.

1. Introduction

Artificial intelligence is increasingly being reshaped by embodied intelligence, where perception and action are continuously coupled through interaction with real-world environments rather than relying solely on computational processes (Liu et al., 2025). Rapid advances in large language models have accelerated progress in the field, while developments in multimodal learning and robotics have further expanded its research scope and technological potential[1].

The relationship between scientific research and technological innovation has long attracted attention in scientometric studies. Early work demonstrated strong connections between scientific publications and patents through citation flows and knowledge spillovers (Narin et al., 1997). Subsequent studies extended this line of inquiry by tracing citations from patents to scientific literature, providing further evidence of science–technology coupling across institutional boundaries (Glänzel & Meyer, 2003). Building on these findings, research on technology opportunity identification gradually shifted toward data-driven methods, with the

aim of detecting emerging technological trajectories before they developed into dominant technological paradigms (Porter & Detampel, 1995; Yoon & Park, 2004)[2-5].

But most existing studies still rely on static topic modeling or single-source analysis, so they do not show well how the relationship between the two domains keeps changing. In embodied intelligence, this weakness looks sharper because publications and patents follow different evolutionary trajectories. Cross-domain temporal dynamics also get too little attention. Then it becomes hard to judge the pattern. Are the two sides moving closer? Is one side pulling away? Or are they really changing together through co-evolution?

To address these limitations, we develop a Topic Evolution and Technology Opportunity (TETO) framework from a dual-source perspective of scientific publications and patents. The framework integrates a Dynamic LDA model with temporal dependence to capture topic evolution over time. Based on cross-domain topic trajectories, it identifies frontier-exploration, industry-driven, and co-evolutionary technology opportunities in embodied intelligence.

This study makes three contributions. On the method side, we extend traditional LDA by adding temporal dependence, and that lets topic evolution be tracked in a dynamic way across time. We also develop the TETO framework and apply it to embodied intelligence, so cross-domain topic change and technology opportunity structures can be examined together. For empirical support, we use 821 scientific publications and 576 patents, and the results show the framework works well.

2. Literature Review

2.1. Embodied Intelligence

In embodied intelligence, the basic view is that intelligence emerges from interaction among brain, body, and environment. In physical contexts, cognition is shaped through continuous perception-action coupling (Liu et al., 2025). Research in this field is usually discussed in three streams, but they do not lean in same direction. One stream is more conceptual and review-based, and it focuses on learning and control, especially reinforcement learning and robotics (Sutton & Barto, 1998; Kober et al., 2013). Another maps knowledge structures through bibliometric methods (Small, 1973; Cobo et al., 2011). Then there is work using topic modeling, such as LDA, to extract latent semantic structures (Blei et al., 2003; Griffiths & Steyvers, 2004)[6-10].

2.2. Topic Evolution Analysis

What topic evolution really asks is pretty direct, how research themes change over time, both in prevalence and in structure. Early studies mostly relied on co-citation and co-word analysis (Small, 1973; Callon et al., 1983). Later, dynamic topic models moved beyond static approaches by bringing in temporal variation (Blei & Lafferty, 2006). After that, longitudinal text analysis gave more fine-grained temporal modeling (Wang & McCallum, 2006; Hall et al., 2008). Even now, most methods still miss cross-system dynamics[11-13].

2.3. Technology Opportunity Identification

Many studies still rely on only publications or only patents, so academia-industry insight remains limited. Technology opportunity identification is really about finding emerging technological directions with innovation potential (Porter & Detampel, 1995). Early work stayed close to knowledge recombination. More recent studies turned to data-driven methods and used text mining together with patent analysis to identify emerging trends (Yoon & Park, 2004; Lee et al., 2009). Topic modeling also improves semantic analysis of technological evolution[14-19].

Even after a lot of research, three limits still remain. Most existing methods are still static, so temporal evolution is hard to trace and cross-system dynamics do not come out clearly. Many studies also stay with one analytical perspective only, which means scientific publications and patents are usually not analyzed together. In embodied intelligence, dynamic topic modeling is also used in only few studies for technology opportunity identification. Because of these gaps, we introduce Dynamic Latent Dirichlet Allocation (LDA) in a dual-source framework of publications and patents, and then develop TETO to analyze topic evolution and identify technology opportunities.

3. Topic Evolution and Technology Opportunity Analysis Based on Improved LDA

Figure 1 shows how the whole framework runs. We first preprocess publication and patent data, then split both sources into time segments. After that, we develop Dynamic LDA and evaluate topic quality. The later part looks at topic evolution, compares the two sources, identifies technology opportunities, and leaves result interpretation for final step.

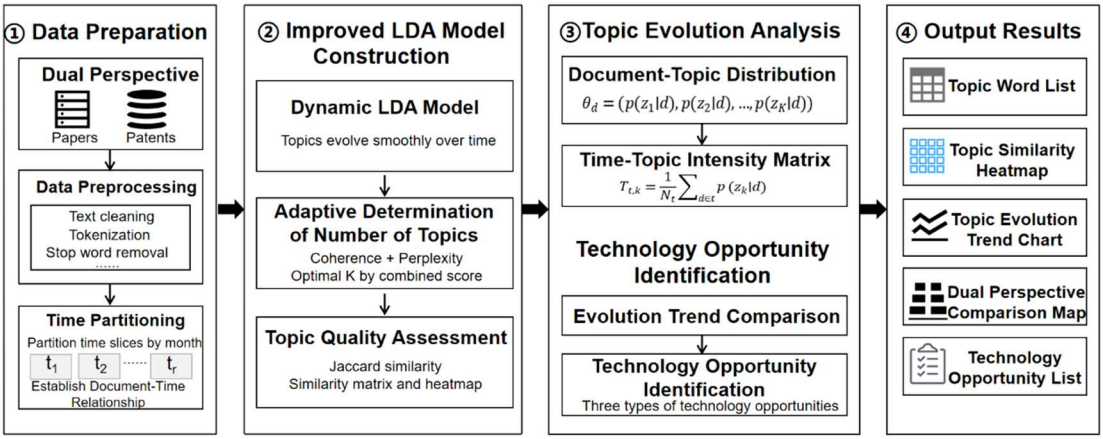


Figure 1. Analytical framework

3.1. Data Preprocessing

For publication data, we use CNKI core journals and CSSCI databases, and this gives 842 articles on embodied intelligence from February 2006 to March 2026. Patent data are retrieved from the China National Intellectual Property Administration (CNIPA), covering 589 patents from July 2020 to February 2026. After data cleaning, 821 publications and 576 patents are retained. Then all texts go through noise filtering and format normalization, and publication dates together with patent dates are unified into monthly time slices for Dynamic LDA modeling.

3.2. Dynamic LDA model

3.2.1. Model Construction

Traditional LDA is not suitable for modeling temporal evolution, because it assumes static topic distributions. Dynamic Topic Models (DTM) can capture variation over time, but computational cost is high, and scalability stays limited when publication and patent data become large-scale.

We build a Dynamic LDA model for this problem by putting temporal information into LDA. Instead of using whole corpus as one block, we split it into monthly time slices by publication dates, so the temporal resolution gets better. Then the topic–word distributions become time-dependent, and topic semantics can change smoothly over time.

If corpus is divided into T time slices, each document d belongs to one time slice td. Then the topic distribution for each document still follows the standard LDA formulation:

$$\theta_d \sim Dirichlet(\alpha)$$

where θ_d denotes the topic distribution of document d over all topics. Unlike standard LDA, we extend the topic–word distribution from a static form to a dynamic one. Let $\beta_k^{(t)}$ denote the word distribution of topic k at time t, and its evolution is defined as:

$$\beta_k^{(t)} \sim \mathcal{N}(\beta_k^{(t-1)}, \sigma^2 I)$$

From this formulation, we know the topic–word distribution changes through a first-order Markov process over time. The topic semantics at current step comes from previous time step, and it shifts in a smooth way. Here, σ2 gives the variance of the evolutionary process.

Based on this, the word generation process in each document is defined as:

$$z_{d,n} \sim Multinomial(\theta_d)$$
$$w_{d,n} \sim Multinomial(\beta_{z_{d,n}}^{(t_d)})$$

where $z_{d,n}$ denotes the topic assignment of the n-th word in document d, and $w_{d,n}$ is the generated word. Standard LDA does not handle it in same way. In this model, word generation follows the topic distribution in the related time slice, so temporal information enters the topic generation process directly.

Then, the joint probability distribution in Dynamic LDA can be written as:

$$p(w, z, \theta, \beta) = \prod_{t=1}^{T} p\left(\beta^{(t)} \mid \beta^{(t-1)}\right) \prod_{d=1}^{M} p\left(\theta_d \mid \alpha\right) \prod_{n=1}^{N_d} p\left(z_{d,n} \mid \theta_d\right) p(w_{d,n} \mid z_{d,n}, \beta^{(t_d)})$$

We want to make topic selection less subjective, so we use an adaptive way based on topic coherence and perplexity. The candidate topic numbers are set as:

$$K \in \{2,3,\cdots,15\}$$

For each K, topic coherence and perplexity are calculated and normalized:

$$Coherence^*(K) = \frac{Coherence(K) - Coherence_{min}}{Coherence_{max} - Coherence_{min}}$$
$$Perplexity^*(K) = 1 - \frac{Perplexity(K) - Perplexity_{min}}{Perplexity_{max} - Perplexity_{min}}$$

A composite score is then constructed:

$$Score(K) = \lambda \cdot Coherence^*(K) + (1 - \lambda) \cdot Perplexity^*(K)$$

The optimal number of topics is selected as:

$$K^* = \arg\max_{K \in \mathcal{K}} \quad Score(K)$$

Next, we also check topic quality with Jaccard similarity, using top keywords taken from each topic:

$$J(A_i, A_j) = \frac{|A_i \cap A_j|}{|A_i \cup A_j|}$$

A higher similarity value means more keyword overlap, so topics become less easy to separate.

3.2.2. Topic Evolution Analysis and Technology Opportunity Identification

From the Dynamic LDA results, we build a topic evolution module in TETO, and this lets us track how topic intensity changes over time.

For each document, the model generates a topic distribution:

$$\theta_d = (p(z_1 \mid d), p(z_2 \mid d), \ldots, p(z_K \mid d))$$

Next, we sort documents into time slices and use them to build a time-topic intensity matrix:

$$T_{t,k} = \frac{1}{N_t} \sum_{d \in t} p\,(z_k \mid d)$$

where $T_{t,k}$ represents the average intensity of topic k at time period t.

We look for technology opportunities by applying Dynamic LDA separately to publication corpus and patent corpus, then we compare topic trajectories and intensity changes between these two domains.

This study does not take highly active topics straight as technology opportunities. Instead, it identifies opportunities from cross-domain asymmetries between scientific research and technological development.

From these evolutionary patterns, the study then divides technology opportunities into three categories.

Frontier-exploration opportunities: topics that grow rapidly in publications but remain weak in patents, indicating strong theoretical exploration potential.

Industry-driven opportunities: topics that grow more rapidly in patents than in publications, reflecting strong industrial demand and application-oriented development.

Co-evolutionary opportunities: topics that exhibit synchronized growth across publications and patents, indicating increasing convergence between scientific research and technological application.

4. Results and Discussion

Using the improved Dynamic LDA model, this part looks at embodied intelligence from both publications and patents. We begin after model estimation by checking if the topic model works well enough, then we trace how topics change over time from academic and industrial sides, and in the end we compare these two sources to see where they diverge, where they come closer, and what technology opportunities may be there.

4.1. Topic Evolution in Publications

4.1.1. Validation of Topic Modeling Results

For choosing the number of topics, we do not want to depend too much on researcher judgment. So we test LDA models with different topic numbers, and we use coherence together with

perplexity to pick the best K. Figure 2 shows that K = 8 gives the most suitable balance between topic clarity and model fit. Because of that, the later Dynamic LDA analysis uses eight topics. We also check topic quality with Jaccard similarity. From Figure 3, most topic pairs stay at low similarity, so the topics are semantically separate and the overall structure is quite stable.

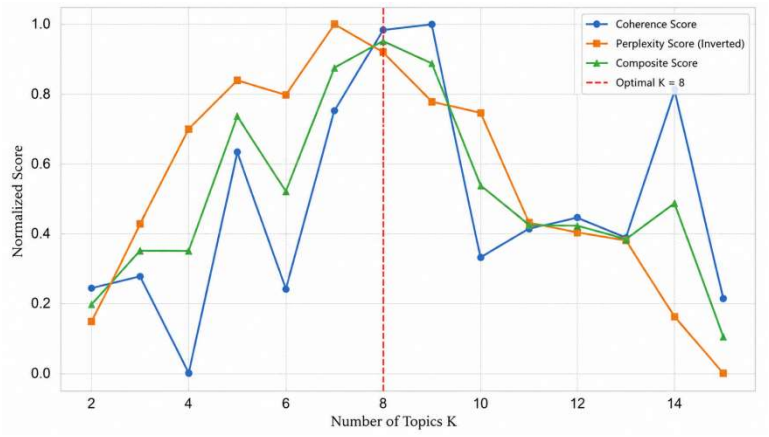


Figure 2. Evaluation of Optimal Topic Number for Publications

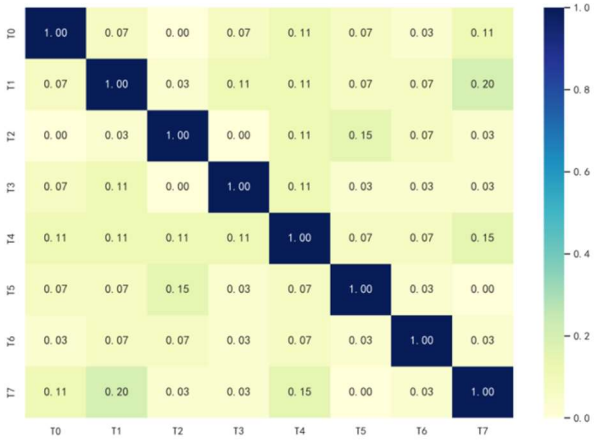


Figure 3. Jaccard Similarity Heatmap of Publication Topics

4.1.2. Topic Identification Results and Topic Evolution Analysis of Publications

Table 1. Publication Topic Identification Results

Topic ID	Topic Name	Keywords
T1	Cognitive Science and Philosophical Foundations	cognition, science, body, consciousness, environment, subject, behavior, mechanism, coupling, experience
T2	Subjective Consciousness and Human-Machine Relations	consciousness, Human-Machine, ethics, interaction, society, collaboration, embodied agent, value, responsibility, trust
T3	Multimodal Perception and Environmental Interaction	multimodal, vision, speech, sensing, environment, fusion, perception, cross-modal, integration, recognition
T4	Robot Control and Motion Execution	control, motion, trajectory, execution, robotic arm, algorithm, optimization, planning, actuation, feedback
T5	Situational Experience and Intelligent Service Applications	context, user, service, healthcare, intelligence, experience, scenario, personalization, interaction, assistance
T6	Industrial Collaboration and Intelligent Applications	intelligence, industry, manufacturing, transformation, system, application, collaboration, production, data, network
T7	Human-Machine Interaction and Task Planning	Human-Machine, task, decision-making, behavior, intelligence, collaboration, control, planning, coordination, learning
T8	Embodied Communication and Educational Applications	education, interactive, media, robotics, learning, communication, immersion, virtual, technology, design

Based on K = 8, we run topic modeling on publication data. The topic names come from top keywords and their semantics, as shown in Table 1. What we find is a shift from basic theory to research that is more about methods and application, and the whole path moves from cognitive and philosophical foundations toward engineering implementation and real-world use.

In Figure 4, later studies already spread into education, healthcare, and service robotics, but at beginning the research was more about cognition and philosophy, then it moved toward multimodal perception, human–machine interaction, and task planning.

By contrast, T4 (Robot Control and Motion Execution) stays quite steady, which suggests this line is already mature and works as foundation. T3 (Multimodal Perception and Environmental Interaction) and T7 (Human-Machine Interaction and Task Planning) rise fast, mainly because large models and embodied agents developed quickly.

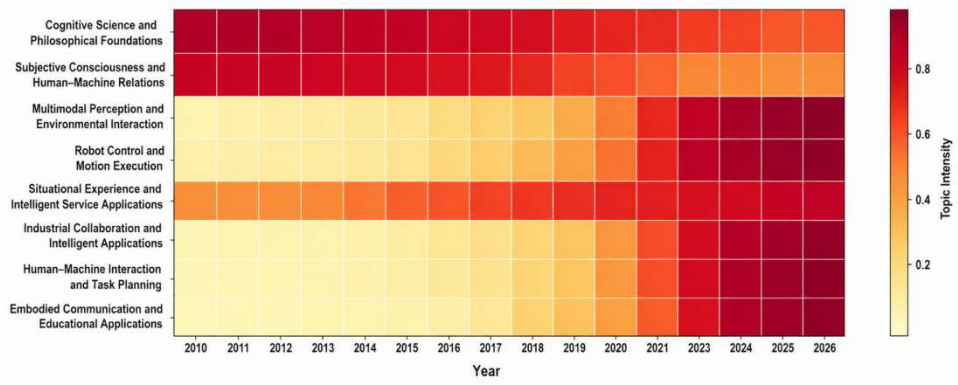


Figure 4. Temporal Evolution Heatmap of Publication Topics

4.2. Topic Evolution Analysis of Patents

Publications tell us more about research side, but patents are closer to industrial implementation and system integration, so patent analysis helps trace the industrialization path of embodied intelligence.

4.2.1. Validation of Patent Topic Modeling Results

For patent data, we use same procedure again. Figure 5 suggests K = 7 gives the best composite score when coherence and perplexity are considered together, so in the end we keep seven topics.

According to Jaccard similarity results in Figure 6, overlap between topics stays low, so the technological boundaries look fairly clear. Patent topics also have stronger local connectivity than publications, and this points to tighter engineering coordination.

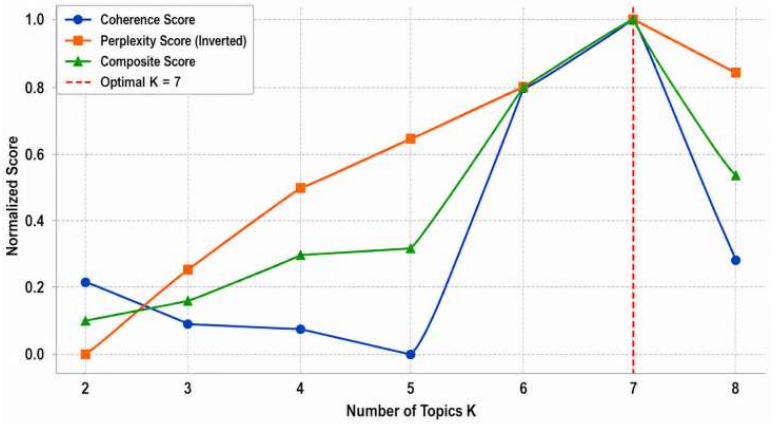


Figure 5. Evaluation of Optimal Topic Number for Patents

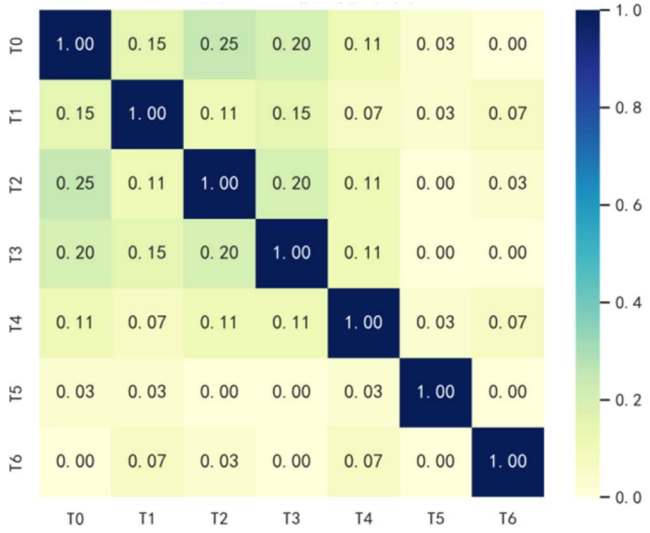


Figure 6. Jaccard Similarity Heatmap of Patent Topics

4.2.2. Patent Topic Identification Results and Topic Evolution Analysis

With K = 7, Table 2 shows what the patent topic model gives. The pattern is mostly on engineering side, and the main attention stays on motion control, system integration, and hardware optimization.

Table 2. Patent Topic Identification Results

Topic ID	Topic Name	Keywords
T1	Motion Control and Actuator Systems	control, drive, motor, joint, robotic arm, servo, actuation, motion, power, stability
T2	Trajectory Optimization and Motion Training	path, trajectory, navigation, obstacle avoidance, motion, training, learning, planning, mobility, localization
T3	Flexible Sensing and Tactile Detection	sensing, tactile, pressure, detection, feedback, perception, flexible, electronic skin, interaction, acquisition
T4	Task Planning and Decision Control	task, decision-making, planning, scheduling, control, coordination, autonomy, execution, strategy, allocation
T5	Dialogue Memory and Behavior Policy Generation	dialogue, memory, semantics, generation, strategy, intelligent agent, large model, interaction, behavior, reasoning
T6	Product Design and Interface Representation	appearance, interface, display, structure, terminal, layout, interaction, visual, design, screen
T7	Intelligent System Integration and Hardware Optimization	system, hardware, module, integration, platform, terminal, chip, device, embedded system, optimization

Figure 7 shows the early stage first, where patents stay mainly around motion control and trajectory optimization. After 2021, the direction changes fast, and topics linked to task planning, decision control, and behavior policy generation rise quickly.

Meanwhile, flexible sensing and system integration keep growing in a steady way, which seems to show more attention going to perception and optimization at system level. Overall, patent focus moves out of basic control and toward intelligent decision-making with more integrated applications.

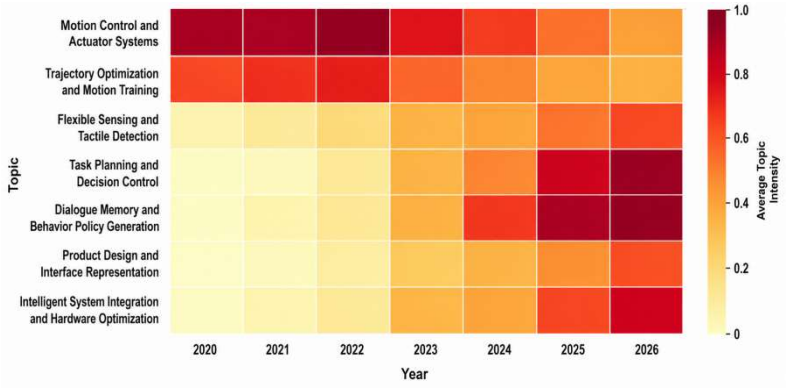


Figure 7. Temporal Evolution Heatmap of Patent Topics

4.3. Dual-source Comparative Analysis of Publications and Patents

4.3.1. Structural and Evolutionary Differences between Publication and Patent Topics

To further uncover the differences between academic research and industrial applications in embodied intelligence, we construct cross-source topic alignment relationships based on semantic similarity between publication and patent keywords, as illustrated in Figure 8. On this basis, we conduct a comparative analysis of topic structures and their temporal evolution within a dual-source framework.

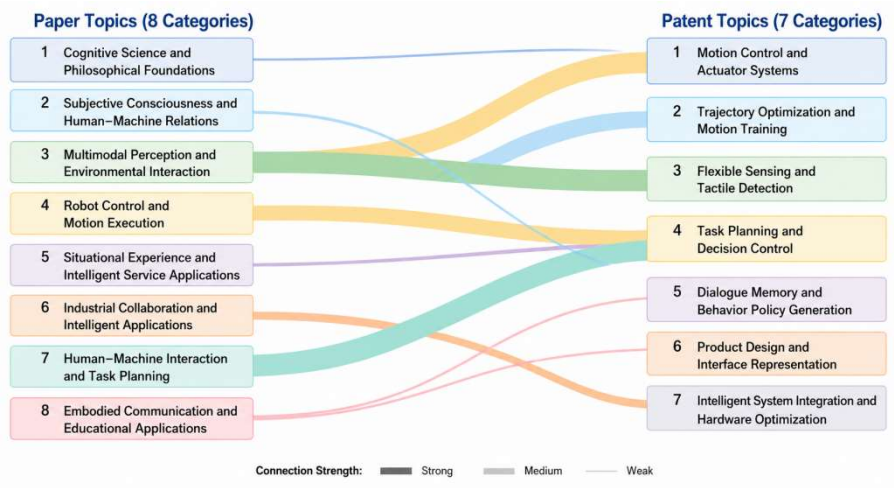


Figure 8. Publication-Patent Topic Mapping

Put side by side, the two knowledge sources are clearly not doing same thing. Publication-based research stays more with cognitive mechanisms, Human-Machine interaction, multimodal perception, and application scenarios, so it asks how intelligence emerges and how systems understand environments. Patents lean much more to engineering work like motion control, actuators, system integration, and hardware optimization, and there the focus is more on system stability, deployability, and implementation efficiency.

That split shows two different innovation logics in dual-source system. Publications mainly serve exploration and theory building, while patents are driven more by application demand and put more weight on engineering implementation and commercialization.

A clear asymmetry appears in topics pushed more by academia. In publications, Cognitive Science stays highly active, and the same is true for Philosophical Foundations and Subjective Consciousness, also for Human-Machine Relations. Patents give these areas much less space, which suggests they are still at theoretical exploration stage, with only limited movement into engineering.

By contrast, industry-driven topics grow faster in patents than in publications. "Motion Control and Actuator Systems" and "Intelligent System Integration and Hardware Optimization" both reflect strong industrial demand for robotics control, system stability, and hardware coordination, so engineering implementation moves ahead quickly.

Some topics, though, are rising in both sources at same time. "Multimodal Perception and Environmental Interaction" is one case. "Human-Machine Interaction and Task Planning" shows same pattern, and "Flexible Sensing and Tactile Technologies" is growing too. In these areas, academic research focuses more on perception mechanisms, multimodal fusion, and decision-making, but patents stay closer to system deployment and implementation.

In multimodal perception and Human-Machine interaction, we usually see topic in publications first, and patent activity comes after. So if we read dual-source system by time, the lag is quite clear. Academic research is normally about 1-3 years ahead of industrial deployment.

Motion control is not same. Because it is more engineering-oriented and driven by demand, patent growth there often appears earlier than academic publication, or it just grows faster. At the same time, dual-source system is moving from control-centric technologies toward intelligent systems that first perceive, then decide, and finally act.

4.3.2. Technology Opportunity Identification

Here, technology opportunities do not mean a topic is simply frequent, and being newly emerging is also not enough. From the dual-source topic evolution results, we look instead at structural asymmetry between academic research and industrial development, or at moments when publications and patents speed up together. When patterns like these show up, the direction is usually still changing fast, and it has not yet reached equilibrium between scientific exploration and industrial deployment.

As summarized in Table 3, technology opportunities in embodied intelligence can be categorized into three types: frontier-exploration-oriented, industry-driven, and co-evolutionary opportunities.

Frontier-exploration-oriented opportunities are primarily concentrated in cognitive science, human agency, and embodied interaction. These topics are highly active in publications but show limited patent activity, indicating that they remain at the stage of theoretical investigation and cognitive mechanism modeling, with substantial room for future technological translation.

Industry-driven opportunities are characterized by stronger growth in patents than in publications. Topics such as motion control, trajectory planning, actuator systems, and system integration reflect strong industrial demand for engineering implementation. These areas are closely linked to robotics control, hardware coordination, and system reliability, and are therefore rapidly advancing toward commercialization.

Co-evolutionary opportunities exhibit simultaneous growth in both publications and patents. Topics such as multimodal perception, Human-Machine interaction, task planning, and flexible sensing suggest that these directions have entered a stage of joint academic-industrial development. In particular, multimodal fusion, autonomous decision-making, and complex environment interaction are emerging as core technological trajectories in embodied intelligence.

Looking at the two sources together, embodied intelligence is going through a continuous evolution instead of changing in one jump. It starts from theoretical exploration, then moves into engineering development, and later reaches industrial deployment; academic research mainly pushes foundational innovation first, while industrial actors turn mature knowledge into application-oriented systems, which makes the pattern sequential but now more and more coupled.

Table 3. Classification of technology opportunities

Opportunity type	Representative Topics
Frontier-Exploration Oriented	Cognitive Science and Philosophical Foundations; Subjective Consciousness and Human-Machine Relations; Situational Experience and Intelligent Service Applications; Embodied Communication and Educational Applications
Industry-Driven	Motion Control and Actuator Systems; Trajectory Optimization and Motion Training; Intelligent System Integration and Hardware Optimization; Product Design and Interface Representation
Co-Evolutionary	Multimodal Perception and Environmental Interaction; Human-Machine Interaction and Task Planning; Task Planning and Decision Control; Flexible Sensing and Tactile Detection; Dialogue Memory and Behavior Policy Generation

4.4. Implications and Recommendations

From the dual-source topic evolution and technology opportunity analysis, embodied intelligence has already moved away from early theoretical exploration and gone further into technological development and application expansion. But different topic types do not follow same path in academic and industrial domains, so policy choices and research strategies need to be differentiated.

For frontier-exploration topics like cognitive science and human-machine relations, I think sustained investment in foundational research is still necessary. These topics stay very active in publications, but their presence in patents is limited, which suggests practical deployment is still quite far away. Future work should pay more attention to embodied cognition mechanisms, and it should keep studying autonomous learning when system faces complex environment adaptation. At same time, interdisciplinary integration with cognitive science, psychology, and even neuroscience is needed. Open platforms and shared datasets should also be built, then experimental validation in education and elderly care can be supported better, and long-term technological accumulation may continue.

For industry-driven topics such as motion control, trajectory planning, and system integration, emphasis should be placed on engineering capability enhancement and industrial coordination. Patent growth is now much faster than paper output, industrial demand is shaping the direction of technology more and more. Because of this, future work should push motion control to higher precision. Planning has to stay dynamic in complex environments, and embodied agents need better system stability too. Joint work on core hardware, such as sensors, chips, plus actuators, should be strengthened, with standardized interfaces and modular architectures to lower integration difficulty and deployment cost.

For co-evolutionary topics such as multimodal perception, Human-Machine interaction, and task planning, greater emphasis should be placed on academia-industry collaboration. These topics exhibit sustained growth in both publications and patents, indicating a transition toward joint scientific and engineering advancement. Future work should establish integrated monitoring mechanisms linking publication and patent dynamics to track technological evolution in real time. Collaborative R&D between universities, research institutes, and industry should be strengthened, particularly in areas such as large-model-driven behavior generation, autonomous decision-making, and natural interaction systems. Continuous accumulation of real-world data and feedback from deployment scenarios will be essential to improve system robustness and practical performance in complex environments.

5. Conclusion

Using the TETO framework developed in this study, we examine embodied intelligence from two sources together, scientific publications and patents, then use that view to trace its evolutionary pattern. The results suggest a structural shift in the field. Earlier work leaned more

toward cognitive exploration, while now the center is moving toward technological implementation and industrial application. On academic side, attention stays mainly on cognitive mechanisms plus multimodal interaction. Industry is different. It puts more weight on motion control, and on system integration, with applications that are more engineering-oriented. From the dual-source comparison, three technology opportunities are identified. One is frontier-exploration-oriented. Another is industry-driven, and the third is co-evolutionary opportunities. Within these, multimodal perception and Human-Machine interaction stand out, task planning shows the same co-evolutionary pattern and suggests academic inquiry is converging more with industrial deployment.

Earlier studies mostly rely on one source and use static topic modeling approaches (Small, 1973; Blei et al., 2003). This study goes another way. Publications and patents are analyzed together inside one unified dynamic topic modeling framework, so we get a dual-source perspective instead of a single-source one. With that framework, technological evolution and opportunity structures in embodied intelligence may be described more fully, and the results can offer useful implications for academia-industry collaboration and technology foresight.

There are still some limitations. Embodied intelligence is developing very fast, current dataset may not fully capture some emerging research directions or the latest technological advances. Future research can bring in more diverse data sources and keep updating the analysis framework, so the findings may become more complete and more timely.

References

[1] Honores-Marín, G., Zolfani, S. H., & Chipulu, M. (2026). From patents to predictive analytics: Leveraging R-GCNs for technological opportunity discovery in converging industries. *Technological Forecasting and Social Change*, 227, 124613.

[2] Gao, X., & Rai, V. (2023). Knowledge acquisition and innovation quality: The moderating role of geographical characteristics of technology. *Technovation*, 125, 102766.

[3] Liu, H., Guo, D., & Cangelosi, A. (2025). Embodied intelligence: A synergy of morphology, action, perception and learning. *ACM Computing Surveys*, 57(7), 1–36.

[4] Liu, H., Guo, D., & Huang, K. (2025). Learning for embodiment and embodiment for learning. *Nature Reviews Electrical Engineering*, 2, 651–653.

[5] Sutton, R. S., & Barto, A. G. (1998). Reinforcement learning: An introduction. *IEEE Transactions on Neural Networks*, 9(5), 1054.

[6] Kober, J., Bagnell, J. A., & Peters, J. (2013). Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11), 1238–1274.

[7] Small, H. (1973). Co-citation in the scientific literature: A new measure of the relationship between two documents. *Journal of the American Society for Information Science*, 24(4), 265–269.

[8] Cobo, M. J., López-Herrera, A. G., Herrera-Viedma, E., & Herrera, F. (2011). Science mapping software tools: Review, analysis, and cooperative study among tools. *Journal of the American Society for Information Science and Technology*, 62(7), 1382–1402.

[9] Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3, 993–1022.

[10] Griffiths, T. L., & Steyvers, M. (2004). Finding scientific topics. *Proceedings of the National Academy of Sciences*, 101(Suppl. 1), 5228–5235.

[11] Wang, X., & McCallum, A. (2006). Topics over time: A non-Markov continuous-time model of topical trends. In *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (pp. 424–433).

[12] Callon, M., Courtial, J. P., Turner, W. A., et al. (1983). From translations to problematic networks: An introduction to co-word analysis. *Social Science Information*, 22(2), 191–235.

[13] Hall, D., Jurafsky, D., & Manning, C. D. (2008). Studying the history of ideas using topic models. In *Proceedings of EMNLP* (pp. 363–371).

[14] Wang, X., McCallum, A., & Wei, X. (2007). Topical n-grams: Phrase and topic discovery, with an application to information retrieval. In *Proceedings of ICDM*.

[15] Porter, A. L., & Detampel, M. J. (1995). Technology opportunities analysis. *Technological Forecasting and Social Change*, 49(3), 237–255.

[16] Yoon, B., & Park, Y. (2004). A text-mining-based patent network: Analytical tool for high-technology trend. *Journal of High Technology Management Research*, 15(1), 37–50.

[17] Lee, S., Yoon, B., & Park, Y. (2009). An approach to discovering new technology opportunities: Keyword-based patent map approach. *Technovation*, 29(6–7), 481–497.

[18] Narin, F., Hamilton, K. S., & Olivastro, D. (1997). The increasing linkage between U.S. technology and public science. *Research Policy*, 26(3), 317–330.

[19] Glänzel, W., & Meyer, M. (2003). Patents cited in the scientific literature: An exploratory study of reverse citation relations. *Scientometrics*, 58(2), 415–428.