ISSN 2409-2665
Journal of Logistics, Informatics and Service Science
Vol. 12 (2025) No. 11, pp. 205-218
DOI:10.33168/JLISS.2025.1114

A Decade of Deep Learning in Engineering: Latent Dirichlet Allocation (LDA) Topic Modeling and Roadmapping

Daesoo Choi

Department of Convergence Engineering, Joongbu University, South Korea
dschoi@joongbu.ac.kr

Abstract. This study analyzes the evolution of deep learning research topics in engineering by applying Latent Dirichlet Allocation (LDA) topic modeling on 2,000 SCIE and SCOPUS-indexed article titles from 2015 to 2024. The research identifies three dominant topics in each of the two-time windows, indicating a shift from foundational methodologies to applied system integration. The results are interpreted through a deep learning technology roadmap, highlighting both thematic continuity and the direction of innovation. This study is novel in its integration of LDA-based trend analysis with the construction of strategic roadmaps specific to engineering domains. Although the use of article titles provides structured insights, it may limit semantic depth compared to full-text analysis. Nevertheless, this research provides structured evidence to support strategic planning for policymakers and academic communities regarding the future direction of deep learning research.

Keywords: deep learning, topic modelling, LDA (Latent Dirichlet Allocation), research trends, technology roadmap, engineering applications, topic evolution, strategic planning

1. Introduction

In recent years, deep learning has emerged as a transformative technology across a wide range of engineering domains, including computer vision, signal processing, and autonomous control systems. This significant growth in research output has resulted in a vast body of academic literature, reflecting rapid advancements in algorithms, architectures, and applications. As the field continues to evolve, understanding the longitudinal trends in deep learning research has become crucial for identifying emerging technologies and forecasting future directions.

To address this need, topic modeling techniques such as Latent Dirichlet Allocation (LDA) have emerged as a dominant approach to uncovering latent semantic structures within large-scale textual data (Blei et al., 2003). LDA enables researchers to classify large corpora into coherent thematic clusters, facilitating the analysis of topic evolution across time (Griffiths & Steyvers, 2004). While bibliometric and keyword-based analyses have previously been employed to summarize deep learning research trends (Alghamdi & Alfalqi, 2015), these methods often fall short of capturing the nuanced changes in research focus across periods.

In the field of engineering, where innovation cycles are relatively fast and interdisciplinary applications are expanding, longitudinal topic modeling offers a powerful approach to monitoring the thematic progression of scholarly output. However, few studies have systematically applied LDA modeling to deep learning research titles across distinct time windows, particularly within engineering domains. Most existing works focus on general AI research or broader computer science contexts, often neglecting engineering-specific contributions that feature unique methodological adaptations and application demands. This oversight can lead to gaps in strategic foresight, as policymakers and educators may lack targeted insights into how deep learning is transforming engineering practices. Moreover, the integration of LDA-based results with technology roadmapping remains underexplored despite its potential to contextualize thematic shifts within broader technological developments (Wang and Blei, 2011).

This study aims to fill that gap by conducting a cross-period topic modeling analysis of deep learning research in engineering from 2015 to 2024. Using 2,000 English-language article titles indexed in SCIE and SCOPUS and collected via the Korea Education and Research Information Service (KERIS), we performed LDA modeling on two subsets: 2015–2019 and 2020–2024. The preprocessing pipeline included tokenization, normalization, and removal of stop words to ensure high-quality topic extraction. Additionally, we developed a deep learning technology roadmap to support the interpretation of thematic trends.

Preliminary LDA results indicate that earlier research (2015–2019) focused on hybrid classification approaches, visual data processing, and neural network-based models. In contrast, recent research (2020–2024) demonstrates a shift toward reinforcement learning, multi-input multi-output (MIMO) systems, automated imaging, and framework integration. These findings are supported by improvements in coherence scores and a reduction in perplexity, indicating more semantically distinct topic distributions in the latter period.

By combining topic modeling and roadmap analysis, this study not only reveals how engineering-oriented deep learning research has transformed over the past decade but also provides insights into future research trajectories. These results can help researchers, policymakers, and technologists identify key areas for growth and potential innovation pathways.

2. Theoretical Background and Related Work

2.1. Deep Learning in Engineering Research

Deep learning, a subfield of machine learning based on artificial neural networks, has become a cornerstone of modern engineering research. Its capacity to model complex nonlinear relationships and learn hierarchical feature representations has significantly advanced applications such as image classification, signal processing, fault detection, and autonomous control systems (Goodfellow et al., 2016). In particular, convolutional neural networks (CNNs) and recurrent neural networks (RNNs) have been extensively utilized in engineering domains that require spatial or temporal data interpretation (LeCun et al., 2015).

Since the mid-2010s, the engineering community has witnessed exponential growth in deep learning–driven research outputs (Wang et al., 2016). The increasing availability of large datasets, high-performance computing resources, and open-source frameworks such as TensorFlow and PyTorch has facilitated this. However, as the field matures, the diversity of research topics has also expanded—from low-level algorithmic innovations to domain-specific applications and hybrid modeling frameworks (Chen et al., 2020).

Understanding how these topics have evolved can provide structured evidence to support strategic planning into the maturity, saturation, or emergence of research trends. This is particularly crucial for identifying strategic areas for future exploration, allocating funding, and developing curriculum.

2.2. Topic Modeling for Trend Analysis

Latent Dirichlet Allocation (LDA), introduced by Blei et al. (2003), is one of the most widely used probabilistic topic modeling techniques. It assumes that documents are mixtures of topics, where a distribution over words represents each topic. This method has proven effective in capturing hidden thematic structures within large-scale corpora, especially in scientific and technical domains.

In recent years, LDA has been increasingly applied in bibliometric and scientometric studies to analyze research trends across disciplines. For example, Griffiths and Steyvers (2004) demonstrated the utility of LDA in extracting meaningful scientific topics from academic articles. More specific to engineering and computer science, Alghamdi and Alfalqi (2015) provided a comprehensive survey of LDA-based topic modeling in text mining applications.

Beyond trend identification, some studies have incorporated LDA into research forecasting frameworks. Wang and Blei (2011) extended the model to collaborative topic modeling, aiming to improve recommendation systems for academic literature and demonstrating the model's adaptability to specialized tasks. However, despite these advances, few studies have integrated LDA topic modeling with technology roadmapping—a methodology that can visually organize the evolution of technologies across time frames.

2.3. Research Gap and Novelty

Although prior studies have examined the broad trajectory of artificial intelligence and deep learning research (Alghamdi & Alfalqi, 2015; Chen et al., 2020), they tend to treat the field as a monolithic entity. Relatively few have conducted granular, engineering-specific topic modeling, particularly with a comparative temporal design. Moreover, existing work often relies on abstracts or full texts, this study leverages article titles as a focused, high-density textual unit that captures essential research themes and has shown empirical effectiveness in prior topic modeling literature (Zhao et al., 2012).

This study distinguishes itself by applying LDA to two temporally segmented datasets (2015–2019 and 2020–2024), with a specific focus on engineering-oriented deep learning research. The resulting topic structures are further contextualized through the construction of a deep learning technology roadmap, which has not been extensively explored in previous literature.

To our knowledge, no previous studies have systematically integrated LDA topic modeling with technology roadmapping in the context of engineering research. This combination enables the interpretation of latent topics within a structured temporal framework, offering actionable insights into research trajectories, innovation maturity, and strategic development priorities.

While previous studies have examined AI and deep learning trends at a general level, their abstraction limits actionable insights in engineering-specific subfields such as intelligent sensing, embedded control, and industrial automation. Without targeted analysis, technology roadmaps and research policies may be misaligned with sectoral innovation patterns.

3. Methodology

3.1. Research Design

To analyze thematic trends and transitions in engineering-focused deep learning research, this study employed a mixed-method approach combining Latent Dirichlet Allocation (LDA) topic modeling and deep learning technology roadmapping. The research was conducted in three major stages: (1) data collection, (2) data preprocessing and LDA modeling, and (3) integration with a deep learning technology roadmap for interpretive analysis.

The goal of this methodology is to uncover latent topic structures in two distinct time windows—2015–2019 and 2020–2024—and to interpret these findings within the context of technological progress and research innovation in engineering domains.

3.2. Data Collection

We compiled a dataset of 2,000 English-language academic paper titles from SCIE and SCOPUS-indexed journals using the Korea Education and Research Information Service (KERIS) academic search infrastructure. Titles were selected based on the following inclusion criteria:

(1) publication in peer-reviewed journals categorized under engineering or computer science;

(2) appearance of domain-relevant keywords such as "deep learning," "neural network," "convolutional," "reinforcement learning," or "autoencoder"; and

(3) publication years strictly within either 2015–2019 or 2020–2024.

We excluded non-English articles, conference proceedings, review-only special issues, and titles lacking domain specificity. Each time window (2015–2019 and 2020–2024) contributed 1,000 titles, yielding a balanced dataset of 2,000 documents.

The adequacy of this sample size is supported by prior LDA-based studies that commonly employ 500 to 2,000 short texts per temporal segment to ensure sufficient topic diversity while preserving semantic coherence and model interpretability. Since article titles are densely informative and structurally consistent across journals, they serve as a compact yet reliable proxy for thematic extraction. Moreover, this sample size strikes a balance between corpus richness and computational efficiency, enabling robust longitudinal comparisons of evolving research trends in engineering-focused deep learning studies.

Titles were selected as the primary unit of analysis due to their consistent structure and high density of thematic keywords, which have been shown to perform effectively in prior topic modeling applications (Zhao et al., 2012).

3.3. Data Preprocessing

Before topic modeling, the textual data were subjected to standard natural language preprocessing procedures to enhance topic coherence and minimize lexical noise during the modeling process. The following steps were performed:

Tokenization: Titles were split into individual word tokens.

Lowercasing: All tokens were converted to lowercase for consistency.

Stopword Removal: Common non-informative words (e.g., "the", "and") were removed based on standard stopword dictionaries (Bird et al., 2009).

Lemmatization: Words were reduced to their base forms to unify variants (e.g., "images" → "image").

Filtering: Tokens of less than three characters and those appearing in fewer than five documents were excluded to reduce noise.

Minimum token length: 3 characters

Minimum document frequency for token retention: 5 documents

These parameters were selected based on established practices in topic modeling literature and confirmed through pilot runs to optimize topic coherence. These steps were implemented using the Natural Language Toolkit (NLTK) and Gensim libraries in Python, both of which are widely used in topic modeling pipelines (Řehůřek & Sojka, 2010).

3.4. LDA Topic Modeling

LDA modeling was performed separately on each period to extract topic distributions and identify dominant research themes. LDA assumes that each document (i.e., title) is generated from a mixture of latent topics, each represented by a probability distribution over words (Blei et al., 2003).

The number of topics was set to 3 for both periods, based on a multi-faceted evaluation that included coherence score, model perplexity, and thematic interpretability. In addition to coherence scoring, we conducted a sensitivity analysis by training LDA models with topic counts ranging from two to six. Models with fewer than three topics tended to oversimplify the corpus, merging distinct research directions into overly broad themes. In contrast, models with more than three topics introduced redundant or highly overlapping clusters, which reduced interpretability. Although four- and five-topic models slightly improved perplexity scores, the added complexity did not translate into more apparent topic distinctions. Therefore, the three-topic configuration was selected as it provided the most coherent, non-redundant, and interpretable representation of the research landscape across both periods.

The models were trained using the Mallet implementation of LDA, which has demonstrated improved performance in terms of coherence and convergence (McCallum 2002). Model performance was evaluated using two key metrics:

(1) Coherence Score: Indicates the semantic consistency of the top words in each topic. Higher coherence indicates stronger topic coherence (Röder et al., 2015).

(2) Perplexity: A statistical measure of how well the model predicts a sample. Lower perplexity implies better model fit (Wallach et al., 2009).

The following results were observed:

(1) 2015–2019: Coherence = 0.4987, Perplexity = –7.94

(2) 2020–2024: Coherence = 0.5332, Perplexity = –8.26

These results suggest an increase in topic cohesion and modeling quality in the more recent period.

3.5. Deep Learning Technology Roadmap Analysis

To further interpret the LDA findings in the context of technological progress, we constructed a deep learning technology roadmap covering the same two time periods (2015–2019 and 2020–2024). The roadmap was developed through a structured mapping process that links topic modeling outcomes to concrete technological developments and engineering milestones. The process involved the following steps:

Topic Interpretation: Each LDA-generated topic was analyzed based on its top-ranked keywords and labeled according to its thematic focus (e.g., “hybrid classification,” “reinforcement learning,” “automated systems”).

Domain Categorization: These thematic labels were then matched to specific technology domains in engineering, such as signal processing, MIMO communications, and autonomous control systems, using relevant literature, IEEE standards, and industrial reports.

Temporal Alignment: Identified domains were aligned with corresponding milestones in engineering practice during the respective periods. For instance, the emergence of reinforcement learning in 2020–2024 topics was linked to the growing adoption of deep reinforcement learning in robotic control and smart manufacturing applications.

Integration: This mapping was synthesized into a roadmap structure that visualizes how foundational research themes (2015–2019) evolved into system-level integration and real-world application (2020–2024). The roadmap thus contextualizes abstract topic distributions within the concrete trajectory of technological innovation.

This structured approach enables the replicable integration of LDA-derived topics with engineering technology trends, thereby supporting the interpretive value of longitudinal topic modeling. This integration allowed for a contextualized understanding of how deep learning research in engineering is not only diversifying but also aligning with real-world system development and industrial demand.

4. Results

4.1. Overview of Topic Modeling Results

Latent Dirichlet Allocation (LDA) was applied to two distinct corpora of deep learning–related research article titles from engineering journals, segmented by publication period: 2015–2019 and 2020–2024. Each model was configured to extract three latent topics, selected based on optimal coherence score and interpretability considerations (Griffiths & Steyvers, 2004; Röder et al., 2015). The extracted topics represent clusters of co-occurring terms that define prevalent research themes during each time frame.

Model quality was evaluated using coherence and perplexity metrics. The 2020–2024 model achieved a higher coherence score (0.5332) and a lower perplexity value (–8.26) than the 2015–2019 model (coherence: 0.4987, perplexity: –7.94), indicating increased semantic cohesion and improved model fit in more recent research outputs (Zhao et al., 2012). Table 1 summarizes the topic modeling results, which will be explained in detail in the following paragraphs. Figures 1 and 2 in the appendix are the visualization results.

Table 1. Topic Modeling Results by Period

Period	Topic	Top Keywords	Interpretation
2015–2019	Topic 0	data, approach, networks, classification, visual, feature, analysis, hybrid, image, channel	Hybrid feature extraction and classification using deep learning techniques in engineering applications.
	Topic 1	model, images, sensing, data, network, networks, face, framework, applications	Image-based sensing and pattern recognition in networked systems.
	Topic 2	image, neural, network, survey, networks, approach, data, method, classification, images	General neural network architecture analysis and survey-oriented studies.
2020–2024	Topic 0	image, systems, reinforcement, model, mimo, network, networks, channel, assisted, automatic	Application of reinforcement learning and MIMO communication systems in automated settings.
	Topic 1	approach, image, system, imaging, network, images, sensing, method, analysis, model	Engineering image analysis using diverse deep learning approaches.
	Topic 2	framework, model, data, images, networks, system, method, approach, automated, analysis	Automated integrated frameworks and systems development based on deep learning.

4.2. Topic Distribution: 2015–2019

The LDA model for the 2015–2019 period identified three dominant research themes:

Topic 0: Hybrid Feature Extraction and Classification (Top keywords: data, approach, networks, classification, visual, feature, analysis, hybrid, image, channel)

This topic focuses on hybrid methodologies for feature extraction and classification in engineering applications. The co-occurrence of terms such as "hybrid," "visual," and "channel" indicates the frequent use of combined spatial-spectral features, often employed in remote sensing, object detection, or biomedical signal processing. Notably, hybrid convolutional–handcrafted models were actively explored for tasks such as fault detection in rotating machinery, hyperspectral image classification, and human activity recognition using wearable sensors. These studies helped bridge the gap between classical signal processing and deep learning architectures (LeCu et al., 2015).

Topic 1: Remote Sensing and Pattern Recognition Applications (Top keywords: model, images, approach, sensing, data, network, networks, face, framework, applications)

This cluster focuses on the engineering applications of pattern recognition, particularly in image-based sensing environments. Keywords like "sensing," "framework," and "face" suggest frequent applications in facial recognition systems and environmental monitoring via UAVs or satellite data. For example, deep learning models have been utilized in smart surveillance systems to detect and track individuals, as well as in agricultural engineering for crop health assessment through aerial imaging. The reference to "networked systems" highlights the role of embedded vision and IoT-based sensor frameworks in early deep learning engineering research (Chen et al., 2020).

Topic 2: Neural Network Architecture and Surveys (Top keywords: image, neural, network, survey, networks, approach, data, method, classification, images)

This topic encompasses foundational research into neural network design and performance evaluation. The frequent appearance of terms likes "survey," "architecture," and "method" indicates a strong presence of comparative studies and architecture reviews during this phase. Examples include comparative analyses of CNNs, RNNs, and DBNs in applications such as energy load forecasting or traffic flow prediction. This period saw a proliferation of papers evaluating model robustness, interpretability, and data efficiency, which contributed to the methodological grounding of deep learning adoption in engineering disciplines (Alghamdi & Alfalqi, 2015; Hoyle et al., 2022).

These topics collectively reflect the early stage of deep learning integration into engineering research, with an emphasis on algorithm development, pattern recognition, and hybrid techniques.

4.3. Topic Distribution: 2020–2024

In contrast, the LDA model for 2020–2024 revealed a notable shift in topic orientation:

Topic 0: Reinforcement Learning and MIMO Systems (Top keywords: image, systems, reinforcement, model, mimo, network, networks, and channel, assisted, automatic)

This topic demonstrates a clear shift toward advanced control algorithms and the integration of communication systems. Reinforcement learning (RL), indicated by terms like "reinforcement" and "automatic," is widely used in autonomous navigation (e.g., drones, self-driving vehicles) and industrial robotics for real-time decision-making. The inclusion of "MIMO" and "channel" connects this trend to wireless communication systems—specifically, deep learning-based beamforming and channel estimation in 5G/6G networks. Projects like DeepMIMO and intelligent edge-based vehicular networks demonstrate this confluence of RL and MIMO in real-world engineering deployments (Mnih et al., 2015).

Topic 1: Imaging Systems and Methodological Innovation(Top keywords: approach, image, system, imaging, network, images, sensing, method, analysis, model)

This topic retains the image-centric focus of earlier years but transitions toward more domain-specific innovation. “Imaging” and “sensing” indicate deeper integration into medical engineering (e.g., tumor detection in mammograms, segmentation of brain lesions in MRI) and industrial quality inspection (e.g., printed circuit board defect detection using CNNs). The references to “approach” and “method” also signal methodological diversity, including the use of transformer-based architectures, self-supervised learning, and meta-learning for improving model generalization in image-rich environments (Chen et al., 2020).

Topic 2: Automated Frameworks and Integrated Systems (Top keywords: framework, model, data, images, networks, system, method, approach, automated, analysis)

This cluster captures the maturity of deep learning systems into fully automated pipelines and frameworks. Keywords like “framework,” “automated,” and “system” reflect the deployment of structured ML platforms such as TensorFlow Extended (TFX), MLflow, and Edge Impulse. In engineering contexts, this includes AI-driven predictive maintenance systems in smart factories, AI-integrated building management systems, and scalable sensor fusion platforms for environmental monitoring. The emphasis on automation suggests a paradigm shift from isolated model training to full-stack integration for enhanced operational efficiency and reproducibility (Schmidhuber, 2015).

These topics represent a maturation of the field, indicating a shift from exploratory and foundational research to applied, system-level integration with emerging technologies, such as reinforcement learning and MIMO-enabled communication.

4.4. Comparative Analysis and Thematic Shifts

A comparative summary of topic keywords across the two periods reveals several notable thematic shifts, which are further substantiated by technological milestones in engineering domains:

The decrease in foundational terms such as “neural,” “survey,” and “classification” indicates a saturation of early research focused on basic architecture comparison and performance benchmarking. During the 2015–2019 periods, key developments included the widespread adoption of CNNs and RNNs for image classification and time-series prediction, such as in fault detection systems for mechanical components or energy demand forecasting.

Conversely, the emergence of terms likes “reinforcement,” “mimo,” and “automatic” in the 2020–2024 period reflects the field’s transition toward intelligent decision-making and autonomous control. For instance, the application of deep reinforcement learning in Google DeepMind’s AlphaStar (real-time strategy gaming AI) and OpenAI's robotic hand dexterity project demonstrated the feasibility of learning-based control in complex and uncertain environments. In parallel, deep learning-enhanced MIMO channel estimation models have been actively deployed in the context of 5G and 6G wireless systems, enabling higher spectral efficiency and dynamic resource allocation in smart communication networks.

Finally, the rise of terms such as “framework,” “system,” and “automated” suggests a move from experimental modeling to integrated engineering pipelines. This is evident in the deployment of TensorFlow Extended (TFX) for model production within Google’s AI infrastructure, as well as in automated quality inspection systems utilizing vision-based deep learning in smart manufacturing (e.g., Samsung’s semiconductor inspection lines or Bosch’s predictive maintenance AI platforms). These cases represent a shift from lab-scale prototyping to production-grade implementation, featuring a traceable, modular architecture.

Taken together, these examples demonstrate how deep learning in engineering has evolved from algorithm-level exploration to application-specific integration, supporting the observed thematic transition between the two study periods (Phaal et al., 2004).

5. Technology Roadmap and Discussion

5.1. Visualization of Topic Evolution

The roadmap not only visualizes topic changes but also contextualizes them within engineering milestones and industrial readiness, making the progression of deep learning research both interpretable and strategically informative.

Figure 3 in the appendix presents the deep learning technology roadmap, which synthesizes thematic patterns derived from LDA modeling across two distinct periods (2015–2019 and 2020–2024). This roadmap is structured in three vertical layers—(1) foundational research themes, (2) emerging capabilities, and (3) system-level engineering integration—mapped chronologically from left to right.

In the foundational layer (2015–2019), topics such as hybrid classification and neural network architecture analysis correspond to the roadmap's depiction of early-stage innovation in feature engineering and supervised modeling. These are linked to the widespread use of CNN-based classifiers in quality inspection and predictive maintenance, as well as survey-based efforts that have shaped engineering curriculum and research priorities.

Transitioning into the emerging capability layer (2020–2024), the LDA-identified rise of reinforcement learning and image-based sensing systems aligns with roadmap nodes involving adaptive control algorithms and domain-specific imaging pipelines. These include real-world implementations such as autonomous UAVs using RL for flight planning and deep image segmentation in industrial and biomedical diagnostics.

Finally, in the system integration layer, the third topic cluster from 2020–2024—centered on automated frameworks and scalable deployment—is reflected in the roadmap's emphasis on end-to-end ML pipelines and modular engineering systems. This includes platforms like TFX and Clara, which are designed for full-cycle AI deployment from data ingestion to inference in operational settings.

The arrows and progression shown in Figure 3 thus trace how conceptual research outputs—identified through topic modeling—evolve into embedded engineering capabilities. This linkage provides a time-anchored interpretation of how deep learning has matured within the engineering field, making the roadmap not only a visualization tool but a synthesis of strategic trajectory.

This roadmap approach helps align latent topic structures with technological developments, offering a time-sensitive perspective on how deep learning research is maturing (Zhao et al., 2012).

6. Discussion and Conclusion

This study examined the evolution of deep learning research topics in engineering by applying LDA topic modeling to 2,000 SCIE and SCOPUS-indexed article titles published between 2015 and 2024. The analysis revealed three dominant themes in each period, reflecting distinct stages in the development and application of deep learning technologies.

In the earlier period (2015–2019), research was characterized by a focus on feature extraction, pattern recognition, and architectural analysis, consistent with the foundational phase of deep learning's integration into engineering domains (Goodfellow et al., 2016; LeCun et al., 2015). The presence of terms such as "hybrid," "survey," and "classification" confirms the exploratory nature of these studies.

In contrast, the more recent period (2020–2024) showed a clear transition toward advanced system integration, including reinforcement learning, multi-agent systems, and MIMO-enabled communications. (Mnih et al., 2015; Schmidhuber 2015). For example, deep reinforcement learning has been successfully applied in autonomous robotic control and complex real-time environments such as AlphaStar and OpenAI's robotic manipulation system (Mnih et al., 2015; Akkaya, I., et al., 2019). Similarly, the integration of deep learning with MIMO communication systems has gained traction in 5G/6G signal processing for improving channel estimation and beamforming accuracy (Ye et al., 2020; Huang, et al., 2021).

The rise of terms such as "automated," "assisted," and "framework" reflects the increasing adoption of scalable deep learning platforms for industrial deployment. For instance, Google's TensorFlow

Extended (TFX) (Baylor et al., 2017) and NVIDIA Clara (Yoo, et al., 2021) exemplify real-world engineering frameworks that support modular model deployment, reproducibility, and pipeline automation. These developments underscore a broader engineering trend where AI transitions from experimental tools to integral components in smart manufacturing, autonomous systems, and medical diagnostics (Wuest, et al., 2016; Wuest, et al., 2016).

By incorporating a deep learning technology roadmap, we contextualized these thematic shifts within broader engineering innovation cycles. This dual-layered analysis—combining topic modeling and strategic foresight—enables a more nuanced understanding of how research priorities evolve and align with technological readiness levels (Zhao et al., 2012).

This dual-layered methodology enables more than just retrospective topic analysis; it supports prospective research planning by aligning empirical data trends with long-term technological forecasts.

(1) Contributions and Positioning within the Literature: This study provides an empirical contribution by conducting a comparative LDA analysis on deep learning research in engineering, based on 2,000 SCIE/SCOPUS-indexed article titles segmented into two periods (2015–2019 and 2020–2024). While previous studies (e.g., Chen et al., 2020; Alghamdi & Alfalqi, 2015) have applied LDA to AI-related corpora, most focus on general-purpose or interdisciplinary datasets and often use abstracts or full texts without temporal segmentation. By contrast, our study is novel in its use of title-only corpora—selected for structural consistency and keyword concentration—combined with technology roadmapping to reveal not only topic evolution but also its alignment with real-world engineering systems development. Future research can expand this work by incorporating abstracts or full texts to capture more nuanced semantic relationships, possibly employing transformer-based models for deeper interpretability (Hoyle et al., 2022).

(2) Methodological Limitations: We acknowledge that limiting the analysis to article titles imposes constraints. Titles may lack the contextual richness and nuanced argumentation found in abstracts or full texts. This could result in the omission of implicit subtopics or the oversimplification of complex methodological innovations. However, titles offer practical advantages: they are standardized, consistently available across datasets, and less susceptible to verbosity or noise—thus improving model convergence and topic coherence in short-text LDA.

(3) Implications and Future Research: Our findings offer actionable insights for engineering educators, research strategists, and funding agencies seeking to understand the trajectory of deep learning research.

While article titles provide a concise basis for topic modeling, they limit semantic depth. Future studies could employ abstract or full-text analysis to capture richer contextual information and identify subthemes that are not visible in titles. Citation network analysis may help identify influential research clusters or hidden topic linkages that lexical models overlook. Additionally, transformer-based models like BERTopic could yield more nuanced and coherent topic structures by incorporating contextual embeddings. These approaches would not only validate but also extend or refine the thematic trends and roadmap insights derived in this study.

References

Akkaya, I., et al. (2019). Solving Rubik's cube with a robot hand. arXiv preprint, arXiv:1910.07113.

Alghamdi, R., & Alfalqi, K. (2015). A survey of topic modeling in text mining. International Journal of Advanced Computer Science and Applications, 6(1), 147–153. https://doi.org/10.14569/IJACSA.2015.060121

Baylor, D., et al. (2017). TFX: A TensorFlow-based production-scale machine learning platform. KDD Systems for ML Workshop.

Bird, S., Klein, E., & Loper, E. (2009). Natural language processing with Python. O'Reilly Media.

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent Dirichlet allocation. Journal of Machine Learning Research, 3, 993–1022. https://doi.org/10.1162/jmlr.2003.3.4-5.993

Chen, C., Zhang, J., & Vogeley, K. (2020). The evolution of deep learning research: A bibliometric analysis. IEEE Access, 8, 140530–140544.

Esteva, A., et al. (2017). Dermatologist-level classification of skin cancer with deep neural networks. Nature, 542(7639), 115–118.

Goodfellow, I., Bengio, Y., & Courville, A. (2016). Deep learning. MIT Press.

Griffiths, T. L., & Steyvers, M. (2004). Finding scientific topics. Proceedings of the National Academy of Sciences, 101(Suppl. 1), 5228–5235. https://doi.org/10.1073/pnas.0307752101

Hoyle, A. M., Goel, P., & Resnik, P. (2022). Is automated topic modeling overhyped? A critical review for computational social science. Communications of the ACM, 65(3), 62–71.

Huang, H., et al. (2021). Deep learning for physical-layer 5G and beyond: Challenges and opportunities. IEEE Wireless Communications, 28(2), 6–13.

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. Nature, 521(7553), 436–444. https://doi.org/10.1038/nature14539

McCallum, A. K. (2002). MALLET: A machine learning for language toolkit. http://mallet.cs.umass.edu

Mnih, V., Kavukcuoglu, K., Silver, D., et al. (2015). Human-level control through deep reinforcement learning. Nature, 518(7540), 529–533. https://doi.org/10.1038/nature14236

Phaal, R., Farrukh, C. J., & Probert, D. R. (2004). Technology roadmapping—A planning framework for evolution and revolution. Technological Forecasting and Social Change, 71(1), 5–26.

Řehůřek, R., & Sojka, P. (2010). Software framework for topic modelling with large corpora. In Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks (pp. 45–50).

Röder, M., Both, A., & Hinneburg, A. (2015). Exploring the space of topic coherence measures. In Proceedings of the Eighth ACM International Conference on Web Search and Data Mining (pp. 399–408).

Schmidhuber, J. (2015). Deep learning in neural networks: An overview. Neural Networks, 61, 85–117.

Shengyi, Z., & Fiaidhi, J. (2023). Building and mining a big data platform for traffic data. Journal of Science and Engineering Research, 2(1), 41–50. https://doi.org/10.56828/jser.2023.2.1.5

Wallach, H. M., Mimno, D., & McCallum, A. (2009). Rethinking LDA: Why priors matter. In Advances in Neural Information Processing Systems (Vol. 22, pp. 1973–1981).

Wang, C., & Blei, D. M. (2011). Collaborative topic modeling for recommending scientific articles. In Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (pp. 448–456). https://doi.org/10.1145/2020408.2020480

Wuest, T., et al. (2016). Machine learning in manufacturing: Advantages, challenges, and applications. Journal of Manufacturing Science and Engineering, 138(12).

Ye, H., et al. (2020). Deep learning for massive MIMO signal detection. IEEE Wireless Communications, 27(4), 74–81.

Yoo, Y., et al. (2021). NVIDIA Clara: AI-accelerated healthcare platform. Journal of Digital Imaging, 34, 537–545.

Zhang, S., & Fiaidhi, J. (2023). Building and mining a big data platform for traffic data. Journal of Science and Engineering Research, 2(1), 41–50. https://doi.org/10.56828/jser.2023.2.1.5

Zhao, W. X., Jiang, J., He, J., & Song, Y. (2012). A collaborative approach to topic modeling for short texts. In Proceedings of the 21st ACM International Conference on Information and Knowledge Management (pp. 1123–1132).

Wang, L., Zhao, Z., & Wu, X. (2016). A Deep Learning Approach to the Classification of 3D Models under BIM Environment. International Journal of Control and Automation, 9(7), 179-188, http://dx.doi.org/10.14257/ijca.2016.9.7.17.

Appendix

Appendix A. Topic Modeling Visualization Results

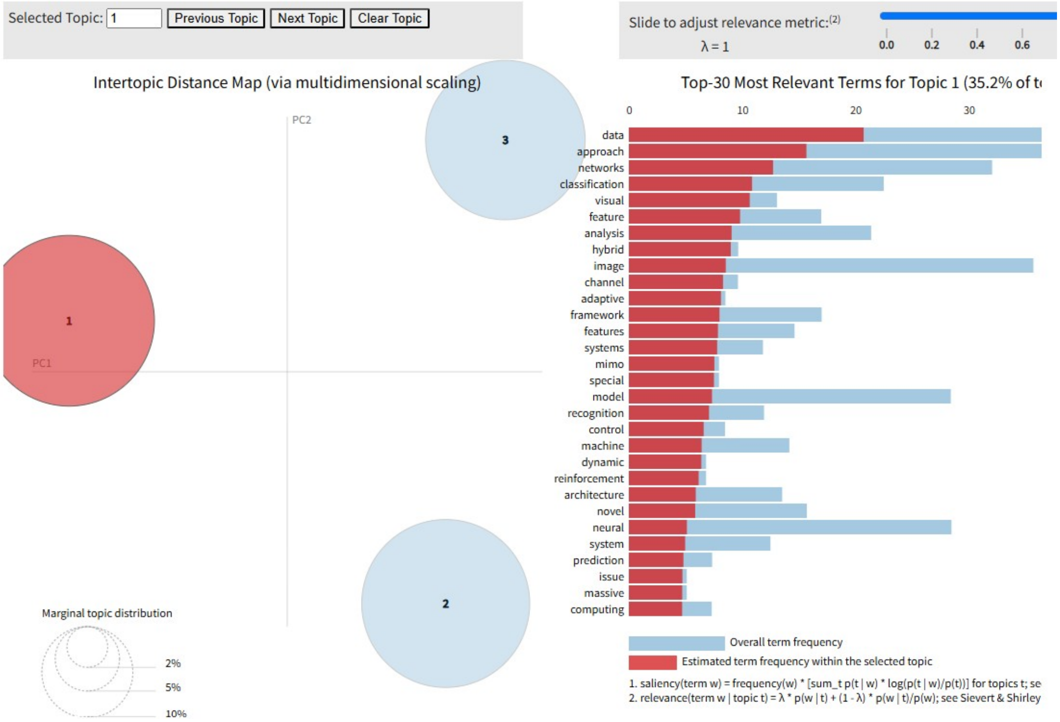

Fig. 1: LDA topic 1 visualization for 2015–2019

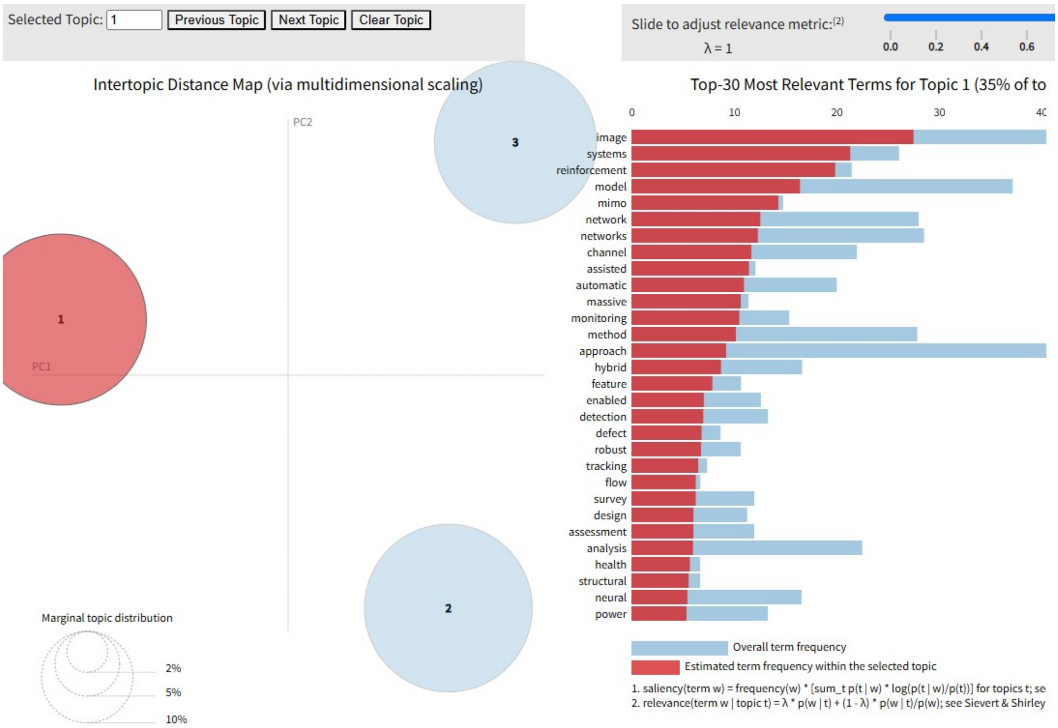

Fig. 2: LDA topic 1 visualization for 2020–2024

Appendix B. Deep Learning Technology Roadmap

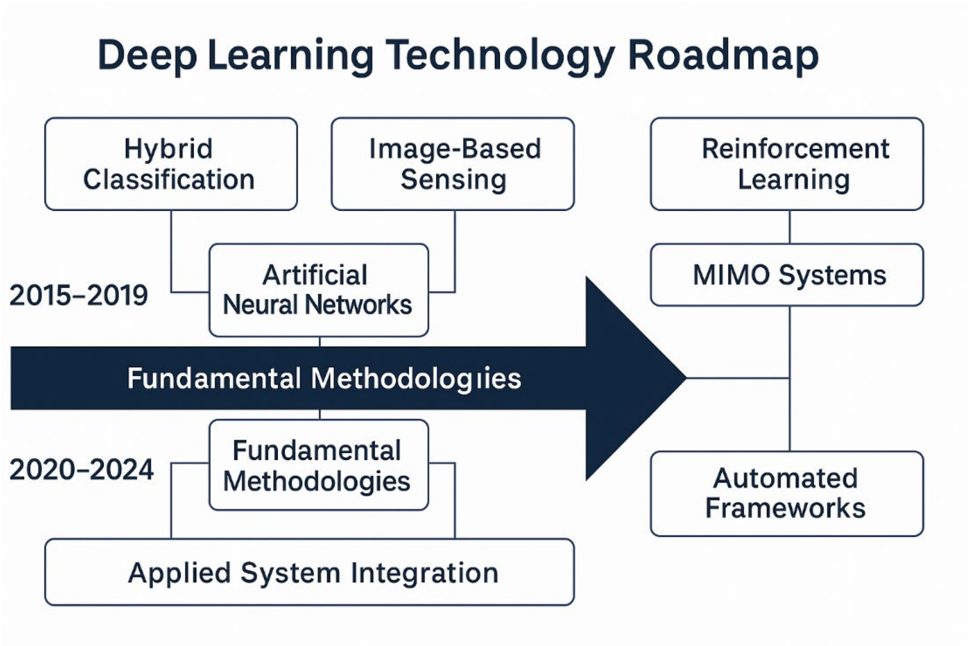

Fig. 3: Deep learning research and technology roadmap (2015–2024)