Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 9, No. 12, 66-80
2025
Publisher: Learning Gate
DOI: 10.55214/2576-8484.v9i12.11273

Disciplinary landscapes of deep learning: Cross-domain insights via LDA topic modeling

Daesoo Choi[1*]
[1]Department of Convergence Engineering, Joongbu University, Goyang, Korea; daesooc100@gmail.com (D.C.).

Abstract: This study employs Latent Dirichlet Allocation (LDA) to analyze research trends in deep learning across engineering, natural sciences, and social sciences from 2020 to 2024. Using a corpus of 3,000 research paper titles, latent thematic structures were extracted to identify the major research directions within each field. The analysis uncovered four prominent topics per domain, revealing clear disciplinary differences in thematic emphasis and levels of methodological maturity. Engineering research predominantly addressed automation technologies, intelligent control systems, and real-time optimization. In contrast, natural science studies focused heavily on medical imaging, computational modeling, and data-driven scientific discovery. Social science research demonstrated an increasing integration of deep learning with ensemble modeling, prediction frameworks, and algorithmic decision processes. By offering a comparative view across disciplines, this study highlights both shared and divergent trajectories in deep learning research. The findings also suggest several future research directions, including the advancement of explainable AI techniques, the incorporation of multimodal data sources, and the development of domain-specific methodological adaptations to improve applicability and interpretability.

Keywords: *Cross-Disciplinary, Deep Learning, LDA, Natural Sciences, Research, Engineering, Social Sciences, Topic Modeling.*

1. Introduction

Deep learning, a transformative subfield within machine learning, has revolutionized the way computers interpret complex data by enabling the automatic extraction of hierarchical features through multi-layered neural network architectures. Unlike conventional machine learning approaches that depend heavily on hand-engineered features and domain-specific preprocessing, deep learning methods excel by learning intricate data representations directly from raw inputs. This end-to-end learning paradigm has led to dramatic improvements in a range of tasks, often surpassing human-level performance in areas such as image classification, speech recognition, and language modeling.

The seminal moment that marked the rise of deep learning can be traced to the 2012 ImageNet Large Scale Visual Recognition Challenge (ILSVRC), where a deep convolutional neural network dramatically outperformed traditional methods. Since then, the success of deep learning has catalyzed its rapid adoption across a wide spectrum of application domains, including but not limited to natural language processing (NLP), computer vision, and speech synthesis. Its ability to model high-dimensional, nonlinear patterns has rendered it particularly valuable in environments characterized by large, unstructured, or noisy datasets.

The academic and industrial impact of deep learning is not confined to the traditional boundaries of computer science. In recent years, its applications have expanded significantly into interdisciplinary and domain-specific fields such as biomedical imaging, financial market forecasting, climate modeling, and smart city development. These fields, which often require the integration of heterogeneous data sources

History: Received: 1 October 2025; Revised: 10 November 2025; Accepted: 14 November 2025; Published: 1 December 2025
* Correspondence: daesooc100@gmail.com

and the uncovering of subtle patterns, benefit greatly from the expressive capacity of deep learning architectures.

Despite this widespread diffusion, much of the existing literature on deep learning remains siloed within individual disciplines. Prior survey studies have predominantly focused on reviewing models and applications within specific scientific domains, often emphasizing methodological innovations or performance benchmarks relevant to that particular field. As a result, there remains a lack of comprehensive, cross-disciplinary analysis that reveals how the research agendas, thematic trends, and topical emphases surrounding deep learning differ or converge across domains.

To address this gap, the present study adopts a data-driven, quantitative approach to investigate how deep learning has shaped research trajectories in three distinct scientific disciplines: engineering, natural sciences, and social sciences. Utilizing Latent Dirichlet Allocation (LDA), an unsupervised probabilistic topic modeling technique, we extract latent thematic structures from a curated corpus of 3,000 research paper titles spanning these three fields. Unlike manual content analysis or narrative reviews, LDA enables the identification of hidden topic distributions that reflect the underlying semantic patterns within large-scale textual data.

Through this lens, we aim to uncover both shared and domain-specific themes in the way deep learning is conceptualized, applied, and discussed in scholarly discourse. For instance, whereas engineering research may emphasize optimization, architecture design, and embedded systems, natural sciences might prioritize data-intensive discovery in areas like genomics or physics simulation, while social sciences could explore interpretability, ethical implications, and human-centered AI systems.

In addition to delineating the topical landscape, this study also provides insights into how deep learning is facilitating interdisciplinary research by bridging methodological and epistemological divides. The comparative perspective afforded by our approach reveals how different academic communities are internalizing and appropriating deep learning technologies, contributing to a more nuanced understanding of its multifaceted impact. By situating our analysis at the intersection of bibliometrics, machine learning, and science mapping, we offer an integrative framework for evaluating the evolving role of deep learning in contemporary research.

Ultimately, this work contributes to the growing body of meta-research that seeks to characterize the diffusion of AI technologies across scientific knowledge systems. It provides empirical evidence that supports strategic decision-making in research funding, curriculum development, and innovation policy, especially in contexts where interdisciplinary integration is both a challenge and an imperative.

2. Theoretical Background and Related Work

Latent Dirichlet Allocation (LDA), first proposed by Blei, Ng, and Jordan, is a seminal probabilistic generative model designed to uncover the hidden thematic structure within large collections of text documents [1]. The core assumption of LDA is that each document in a corpus is composed of multiple latent topics, and each topic is defined as a probability distribution over words. By inferring these distributions, LDA enables the extraction of coherent semantic patterns that are not readily observable through surface-level analysis. As an unsupervised method, it does not rely on labeled data, making it particularly suitable for exploratory data analysis in large and heterogeneous textual corpora.

Over the past two decades, LDA has become a foundational tool in natural language processing (NLP), information retrieval, digital humanities, and bibliometric studies. It has been applied to a wide range of tasks, including document classification, content recommendation, scholarly trend analysis, and sociolinguistic research [2, 3]. A key advantage of LDA is its ability to yield interpretable topic structures that can be visualized and analyzed both qualitatively and quantitatively. Compared to supervised classification models, which require labeled datasets and tend to be tailored to specific tasks, LDA offers a flexible and domain-agnostic framework for discovering latent semantics embedded in text collections.

Numerous enhancements and extensions to the original LDA formulation have been proposed to address limitations such as topic incoherence, lack of temporal dynamics, and inability to model topic correlations. Hierarchical LDA introduces tree-like topic structures that capture multi-level thematic hierarchies [4] while Dynamic Topic Models (DTMs) enable the modeling of topic evolution over time by introducing temporal dependencies between topic distributions [5]. Correlated Topic Models (CTMs) extend the Dirichlet prior with logistic normal distributions to better capture correlations among topics, making them particularly useful in corpora where topics co-occur frequently and meaningfully.

The interpretability of topic models has also been enhanced through the development of advanced visualization tools such as LDAvis, which provides an interactive interface for exploring topic-term relationships and topic prevalence across the corpus [6]. These tools facilitate both expert-driven validation and communication of findings to broader audiences, including those without technical expertise in machine learning.

In terms of domain-specific applications, LDA has proven to be remarkably adaptable. In biomedical research, topic modeling has been used to identify patterns in disease-related publications, gene interaction networks, and drug development literature [7]. Environmental scientists have employed LDA to track climate change discourse across scientific and policy-oriented texts, revealing how narratives around sustainability and global warming have shifted over time [8]. In the field of digital journalism, LDA has facilitated the detection of ideological bias, sentiment framing, and agenda-setting patterns within news articles [9]. These studies highlight the model's versatility and its potential to reveal non-obvious insights in domain-specific contexts.

Despite these successes, relatively few studies have systematically compared topic structures across different academic disciplines using a unified methodological framework. Most existing research remains confined within disciplinary silos, focusing on individual domains without engaging in comparative or integrative analysis. This lack of cross-disciplinary perspective represents a critical gap in the literature, particularly as artificial intelligence and deep learning continue to influence a broad range of scientific and social research areas. Understanding how the discourse surrounding deep learning differs across fields such as engineering, natural sciences, and social sciences is essential for both scholarly synthesis and policy planning.

Recent advancements in topic modeling have begun to incorporate deep learning architectures to overcome the limitations of traditional probabilistic models. Neural topic models, for example, combine the interpretability of LDA with the representational flexibility of neural networks. These models often leverage variational autoencoders (VAEs) or other generative frameworks to learn topic distributions in a continuous latent space, allowing for more nuanced topic semantics. Such approaches offer improved performance on short texts and noisy data while maintaining a degree of interpretability through structured priors.

Moreover, hybrid methods that integrate topic models with graph-based learning techniques or citation network analysis are gaining prominence in bibliometrics. By linking textual content with structural metadata, such as co-authorship, citation frequency, or publication venues, these methods enable a richer contextualization of topic distributions. They are particularly useful for tracking the diffusion of ideas, the emergence of new subfields, and the clustering of research communities around shared methodological or thematic concerns.

In the context of longitudinal and interdisciplinary trend analysis, topic modeling plays a critical role in uncovering macro-level patterns that span across time and disciplinary boundaries. It supports meta-analytical approaches that examine the evolution of research focus, the adoption of methodologies, and the framing of technological innovations like deep learning. As science becomes increasingly interconnected, such tools are indispensable for navigating the complexity of contemporary research ecosystems.

In summary, while LDA and its extensions have been successfully applied across numerous domains, there is a growing need for cross-disciplinary applications that systematically compare how foundational technologies like deep learning are conceptualized and utilized in distinct academic contexts. This study contributes to filling that gap by employing topic modeling as a lens to examine the thematic landscape of deep learning research across engineering, natural sciences, and social sciences, ultimately offering a more integrative view of how this technology is shaping knowledge production across the disciplinary spectrum.

3. Materials and Methods

To investigate the cross-disciplinary landscape of deep learning research, this study conducted a comprehensive topic modeling analysis using Latent Dirichlet Allocation (LDA) across three major scientific domains: engineering, natural sciences, and social sciences. The primary objective was to compare and contrast thematic structures emerging from research discourse within each field. To ensure robustness and relevance, all methodological procedures were meticulously designed, encompassing corpus construction, data preprocessing, model training, evaluation, and interpretability validation.

3.1. Data Collection and Corpus Construction

The data corpus was compiled from peer-reviewed journal articles published between 2020 and 2024, focusing specifically on works that explicitly engaged with deep learning technologies. A total of 3,000 research paper titles were collected, with 1,000 titles assigned to each disciplinary corpus. Titles were retrieved from reputable academic databases such as IEEE Xplore, ScienceDirect, SpringerLink, Web of Science, and Scopus, using targeted keyword queries that included terms such as "deep learning," "neural networks," and "AI applications." Only articles published in English were included to maintain linguistic consistency across the datasets.

To ensure discipline-specific relevance, publications were classified based on the indexing subject category of the journal or conference, as well as manual verification of the research scope. For example, engineering titles included domains such as control systems, robotics, and signal processing; natural sciences encompassed biology, chemistry, and environmental science; and social sciences covered areas like education, psychology, communication studies, and economics.

3.2. Text Preprocessing

Raw title texts were subjected to a standardized preprocessing pipeline to ensure compatibility with LDA modeling requirements. Preprocessing was carried out using a combination of Python libraries, including NLTK and spaCy, chosen for their extensive NLP capabilities and integration with lemmatization tools [10].

The preprocessing steps consisted of the following:

1. Lowercasing: All characters were converted to lowercase to avoid case-based discrepancies.
2. Removal of punctuation and numeric characters: Non-alphabetic symbols were stripped to elimi nate noise from formulaic notations or numbering systems.
3. Tokenization: Titles were split into individual word tokens using whitespace and syntactic boun daries.
4. Stopword removal: Commonly used English stopwords were filtered out using NLTK's built-in stopword list, supplemented by domain-specific additions when appropriate.
5. Lemmatization: Words were reduced to their base or dictionary form to consolidate variants (e. g., "studying," "studies," and "study" were all lemmatized to "study").

Documents with fewer than three valid tokens after preprocessing were excluded from further analysis to ensure sufficient textual substance for topic modeling.

3.3. LDA Model Implementation and Training

For topic extraction, we utilized the Gensim implementation of LDA, which is widely recognized for its computational efficiency and compatibility with large-scale text corpora [4]. To maintain consistency and interpretability across all three disciplinary corpora, the number of topics was initially fixed at four. This decision was informed by prior studies suggesting that lower topic numbers enhance

DOI: 10.55214/2576-8484.v9i12.11273

interpretability in short-text contexts such as titles, and was later verified through coherence score optimization.
Each LDA model was trained independently on its respective disciplinary corpus. The following hyperparameters were employed:

- Number of topics: 4
- Alpha: Set to 'auto' to allow dynamic adjustment of the Dirichlet prior
- Passes: 10
- Random seed: 42, to ensure model reproducibility

Additionally, multiple models with topic numbers ranging from 3 to 10 were trained during the exploratory phase to assess the optimal topic structure. This process involved human-in-the-loop inspection in combination with coherence scoring and perplexity-based evaluation.

3.4. Evaluation Metrics and Model Selection

Model evaluation was conducted using both intrinsic and extrinsic metrics. The primary coherence metric used was the c_v score, which assesses the degree of semantic similarity among the top keywords within each topic. This measure is known to align well with human judgments of topic quality [11]. To complement this, the U_MASS coherence score was also computed to provide a probabilistic estimate of internal consistency. For generalizability assessment, perplexity scores were calculated on held-out validation subsets to estimate how well the model could predict unseen data [12].

To ensure robustness, all three disciplinary LDA models were subjected to the same evaluation pipeline. Final topic models were selected based on the highest average c_v coherence score, with u_mass and perplexity serving as secondary diagnostics.

3.5. Visualization and Topic Interpretation

Interpretability of the resulting topic models was addressed using both quantitative and qualitative methods. Quantitatively, we employed LDAvis, an interactive visualization tool that enables users to explore the relationship between topics and their constituent keywords in two-dimensional space using multidimensional scaling (MDS) [6]. This visualization facilitated the assessment of inter-topic distance, topic dominance, and keyword salience, which proved critical in assigning meaningful labels to each topic.

Qualitatively, manual topic validation was performed by three domain experts, each specializing in one of the target disciplines. These experts reviewed the top-ranked keywords and a sample of associated titles for each topic, providing descriptive labels and assessing thematic coherence. Discrepancies in interpretation were resolved through consensus.

3.6. Domain-Specific Model Structuring

To minimize interpretational confounding due to topic overlap across domains, we adopted a domain-separated modeling strategy. Rather than combining all titles into a single global model, which could obscure field-specific nuances, we trained independent LDA models for each domain. This approach preserved the internal semantics of each corpus while enabling more precise comparisons across domains.

Furthermore, we computed the marginal topic distributions for each corpus to examine topic prevalence and identify dominant research directions. This helped contextualize how deep learning is framed and emphasized differently across disciplines. For example, a topic centered around "model optimization" may be dominant in engineering, while "data ethics and fairness" may surface more prominently in social sciences.

3.7. Methodological Summary

In sum, the materials and methods employed in this study integrate rigorous preprocessing, robust unsupervised modeling, domain-specific validation, and state-of-the-art visualization. By combining LDA-based topic modeling with expert-driven interpretation and cross-domain comparison, we provide a comprehensive methodological framework for analyzing the intellectual structure of deep learning research across multiple scientific disciplines. This approach ensures both analytical depth and interpretive clarity, thereby enabling the subsequent sections to offer empirically grounded insights into the thematic evolution of deep learning in contemporary scholarship [13].

4. Results

This section presents a detailed account of the LDA topic modeling outcomes across the three target domains: engineering, natural sciences, and social sciences. For each discipline, four distinct latent topics were identified based on their most representative keywords. These thematic structures were further interpreted in the context of domain-specific research practices and scientific objectives. To assess the robustness and validity of the generated topics, evaluation metrics such as coherence score and perplexity were employed. These quantitative indicators were complemented by qualitative assessments to ensure interpretability and semantic coherence.

The modeling outcomes reveal that deep learning research manifests with distinctive thematic priorities across the three disciplines, reflecting their epistemological orientations, data practices, and technological maturity. In the engineering domain, topics are strongly centered around system-level implementation and control architecture. For instance, Topic 1 shows frequent use of keywords such as model, control, and monitoring, indicating the integration of deep learning into embedded systems and automation workflows. This thematic focus suggests a discipline primarily concerned with optimizing real-time responsiveness, fault detection, and reliability in intelligent systems. The recurrent presence of terms like network, channel, and tracking further supports this observation, indicating a convergence between communication engineering and reinforcement learning techniques.

In the natural sciences, the topic structures were markedly shaped by data-driven experimentation and imaging technologies. Keywords such as MRI, assessment, tomography, and diagnostic prominently emerged, highlighting the application of deep learning in biomedical image analysis and structured data modeling. This aligns closely with ongoing efforts in computational biology, radiology, and environmental sensing, where AI methods are used to enhance diagnostic accuracy, automate complex measurement pipelines, and facilitate data-intensive discovery. The relatively high coherence scores obtained in this domain underscore the conceptual tightness and maturity of deep learning research in the natural sciences, where model performance is often evaluated within highly standardized experimental protocols.

In contrast, the social sciences exhibited a more diverse and exploratory set of topic patterns. Thematic clusters included ensemble learning methods, socioeconomic forecasting, behavioral analytics, and policy modeling. Representative keywords such as forecasting, ensemble, framework, and algorithm point to a growing interest in integrating computational models into the analysis of complex human systems. The relatively lower coherence scores observed for this domain reflect the epistemic heterogeneity of the social sciences, where qualitative and quantitative methodologies often coexist. This diversity, however, should not be interpreted as a deficiency; rather, it suggests that deep learning is still in the process of being culturally and methodologically assimilated into social science research practices.

Despite these disciplinary distinctions, the topic models also revealed several shared methodological trends. Across all domains, keywords such as model, analysis, approach, and framework were frequently observed. This convergence indicates that while the applications of deep learning differ by field, the underlying computational strategies are becoming increasingly standardized. The widespread adoption of hybrid modeling frameworks, sensor-integrated architectures, and automated image-based systems suggests a cross-disciplinary movement toward modular, reusable deep learning solutions.

4.1. Engineering

The engineering domain revealed highly coherent and application-driven topics, reflecting the field's long-standing emphasis on technological implementation, system optimization, and real-time control. The four dominant themes identified through topic modeling encapsulate key areas where deep learning is being operationalized to enhance system intelligence, autonomy, and adaptability.

- Reinforcement learning and MIMO communication systems: One major topic cluster focused on the integration of reinforcement learning algorithms with multiple-input multiple-output (MIMO) architectures. The presence of keywords such as *mimo*, *channel*, *reinforcement*, and *tracking* indicates a strong research focus on adaptive signal processing, network throughput optimization, and real-time learning in wireless communication environments. This area benefits from deep learning's ability to model dynamic, high-dimensional input spaces and adjust decision policies in uncertain or variable conditions.
- Image-based automation and hybrid neural networks: Another key topic involved the use of convolutional and hybrid neural network models for automated image analysis. Keywords like *image*, *automatic*, *review*, and *optical* suggest that engineering applications are advancing in areas such as industrial inspection, object detection, and computer-aided visual assessment. These technologies are increasingly deployed in manufacturing pipelines, quality assurance systems, and robotic vision, where accuracy, speed, and scalability are paramount.
- Sensor-integrated control and monitoring systems: Deep learning is also being embedded in cyber-physical systems where sensor data must be interpreted and acted upon in real time. Topic terms such as *monitoring*, *control*, *modeling*, and *sensing* indicate a fusion of traditional control theory with neural inference mechanisms. Applications span from smart grids and HVAC systems to autonomous vehicles, where sensor fusion and predictive modeling enable context-aware control under uncertain or rapidly changing conditions.
- Frameworks for robust pattern recognition: The fourth thematic cluster centers around architectural frameworks and algorithmic designs for robust feature extraction and pattern classification. Keywords like *framework*, *features*, *planning*, and *architecture* imply a concerted effort to develop generalizable systems that can be adapted across various engineering domains. This research is particularly relevant for scalable deployment in edge AI and embedded inference applications.

Table 1.
Top 10 keywords extracted from LDA topics for each domain.

Discipline & Topic	Top Keywords	Interpretation
Engineering - Topic 0	MIMO, image, reinforcement, systems, massive, networks, network, channel, tracking, assisted	MIMO systems and RL in communications
Engineering - Topic 1	image, imaging, approach, network, automatic, review, automated, optical, images, analysis	Image processing and automated review systems
Engineering - Topic 2	framework, images, fast, planning, data, detection, automated, architecture, analysis, features	High-speed image processing frameworks
Engineering - Topic 3	model, approach, method, systems, control, channel, modeling, image, monitoring, sensing	Modeling and sensing for control systems

Table 1 summarizes the 10 main keywords for each topic, and Figure 1 in the appendix visualizes the results. The coherence score for the engineering corpus was 0.5607, and the perplexity was measured at -8.3560, indicating a favorable balance between topic clarity and predictive performance. Notably, frequently appearing terms such as *network*, *channel*, and *monitoring* reinforce the inference that deep learning is being strategically positioned within the infrastructure of real-time systems. These findings are consistent with contemporary advances in robotics, autonomous navigation, and smart manufacturing domains [13, 14].

Furthermore, the emergence of terms like *fault*, *vehicle*, and *sensor* reveals an increasing concern with reliability and responsiveness hallmarks of modern engineering systems operating at the edge. The integration of deep learning into these contexts is facilitating more adaptive, intelligent control mechanisms that go beyond rule-based logic, enabling proactive diagnostics, context-sensitive actuation, and continuous system optimization.

Taken together, the engineering domain's topic landscape reflects a mature and technologically aligned adoption of deep learning, emphasizing *functionality*, *performance*, and *operational efficiency*. This suggests not only that deep learning has found a stable foothold within engineering workflows but also that it is actively transforming the design principles and architectures of next-generation engineered systems.

4.2. Natural Sciences

In the natural sciences, the LDA topic modeling revealed four thematically distinct and analytically cohesive topic clusters, each reflecting a strong orientation toward empirical research, quantitative rigor, and domain-specific problem solving. Compared to engineering and social sciences, the topics in this field were markedly characterized by high-dimensional data analysis and precision diagnostics. These results align with the natural sciences' methodological emphasis on reproducibility, data accuracy, and hypothesis-driven inquiry.

- Medical image analysis for cancer detection using MRI and CT: The most prominent topic centers around the application of deep learning techniques in clinical radiology, particularly for oncology-related tasks. Keywords such as *image*, *cancer*, *MRI*, *tomography*, and *automatic* indicate intensive use of convolutional neural networks (CNNs) and transformer-based models in detecting and classifying cancerous tissues from MRI, CT, and PET scans. This research area is critical for advancing early diagnosis, reducing false positives, and enabling non-invasive screening protocols.
- Geospatial and seismic data modeling: Another major topic dealt with the modeling of Earth science data, especially seismic signals and geographic information systems (GIS). The appearance of keywords like *seismic*, *surface*, *mapping*, and *workflow* points to a growing adoption of deep learning for geophysical signal processing, landform recognition, and hazard prediction. These techniques are used to improve the resolution and speed of earthquake detection, map fault lines, and support natural resource exploration.
- Computer-aided diagnosis systems and treatment planning: A third topic focused on systems designed to support or automate clinical decision-making. Keywords such as *assessment*, *dose*, *prostate*, *assisted*, and *study* suggest efforts to refine treatment protocols in radiation therapy, urology, and personalized medicine. Deep learning is applied here to predict patient outcomes, optimize therapeutic regimens, and assist radiologists and oncologists with interpretive tasks.
- Multimodal assessment frameworks: The fourth topic was more methodological, involving the fusion of heterogeneous data types for integrated analysis. Terms such as *data*, *cell*, *framework*, and *models* reflect the use of deep learning to combine imaging, genomic, and biophysical data streams. This multimodal approach is central to systems biology and translational research, where understanding complex interactions requires synthesizing data across scales and modalities.

Table 2.
Top 10 keywords extracted from LDA topics for Natural Science.

Discipline & Topic	Top Keywords	Interpretation
Natural Sciences - Topic 0	image, cancer, model, method, approach, automatic, tomography, framework, analysis, MRI	Medical imaging and cancer detection
Natural Sciences - Topic 1	image, analysis, method, review, mapping, system, workflow, surface, framework, seismic	Seismic and surface data analysis
Natural Sciences - Topic 2	image, MRI, study, assisted, approach, data, dose, prostate, novel, models	MRI-based clinical support studies
Natural Sciences - Topic 3	image, data, model, assisted, assessment, approach, analysis, seismic, cell, system	Scientific data assessment and modeling

Table 2 summarizes the 10 main keywords for each topic, and Figure 2 in the appendix visualizes the results. The coherence score for the natural sciences domain was 0.5868, the highest among all three fields, indicating a particularly strong internal thematic consistency. The associated perplexity score, while not explicitly reported here, was also found to be lower than that of the social sciences corpus, reinforcing the interpretive clarity and semantic tightness of the topics. The recurrence of terms like *MRI*, *assessment*, and *quantitative* reflects the domain's prioritization of precise measurement, methodological standardization, and clinical applicability [15, 16].

These findings are consistent with broader trends in biomedical and environmental sciences, where large-scale structured data such as imaging arrays, time-series signals, and spatial measurements play a foundational role. In biomedical contexts, deep learning accelerates the pace of diagnostic innovation, reducing diagnostic latency and enabling more proactive disease management. In environmental applications, similar models are used to process satellite imagery for land cover classification, track climate change indicators, and monitor biodiversity through remote sensing.

Moreover, the adoption of explainable AI (XAI) principles is gaining traction within this domain, as clinical and scientific end-users increasingly demand interpretability alongside performance. Models must not only be accurate but also justifiable in terms of medical and scientific reasoning. This has led to a rise in hybrid models that combine deep learning with rule-based systems or statistical inference frameworks, particularly in high-stakes applications.

In sum, the topic landscape in natural sciences demonstrates a sophisticated and domain-integrated use of deep learning. These technologies are no longer experimental tools but are being embedded into core scientific workflows, from hypothesis generation and experimental design to diagnostic reporting and environmental monitoring. The methodological rigor and high coherence of the topics suggest that the natural sciences are among the most mature domains in applying deep learning to domain-specific challenges, both in theoretical sophistication and translational impact.

4.3. Social Sciences

The social sciences domain exhibited the most diverse, exploratory, and methodologically pluralistic set of topic structures among the three disciplines analyzed. The LDA results suggest that while deep learning is less entrenched in this field compared to engineering or natural sciences, its presence is nonetheless growing, particularly in areas concerned with prediction, behavioral modeling, and algorithmic governance. The four topics identified reflect this early-stage but expanding adoption of computational methods.

- Social and economic forecasting models: One major topic centered on the use of deep learning to model and predict complex socio-economic phenomena. Keywords such as *forecasting*, *model*, *data*, and *networks* suggest applications in economic trend prediction, public health surveillance, and labor market modeling. The incorporation of deep neural networks allows social scientists to engage with larger, more dynamic datasets, including time-series and spatiotemporal social indicators, which traditional econometric models may not fully capture.
- Image and network-based social data analysis: Another topic revealed a methodological shift toward visual and network-centric data interpretation. Keywords such as *image*, *network*, *design*, *neural*, and *automatic* indicate the use of deep learning for analyzing digital artifacts such as social media images, facial expression data, and online interaction networks. This theme intersects with computational communication and media studies, where neural network models are increasingly applied to understand the diffusion of ideas, visual culture, and online behavior patterns.
- Ensemble and hybrid modeling approaches: A third topic focused on ensemble methods and multi-model integration strategies. Terms like *ensemble*, *hybrid*, *prediction*, and *framework* reflect the field's experimentation with combining deep learning models with traditional statistical techniques such as regression, decision trees, or Bayesian inference. This hybridization trend suggests that researchers in social sciences are seeking to retain interpretability while leveraging the predic

DOI: 10.55214/2576-8484.v9i12.11273

tive power of deep models, particularly in contexts involving uncertainty and heterogeneous dat a sources.

- Algorithmic policy frameworks and behavior modeling: The final topic concerned the role of alg orithms in policy-making and behavior analysis. Keywords including *policy, algorithm, power, fra mework*, and *behavior* point to a growing interest in understanding the societal impact of algorith mic systems, as well as using AI for behavioral forecasting and governance simulation. This the me bridges computational social science with public administration, ethics, and political science, suggesting that deep learning is beginning to inform policy design, regulation, and evaluation.

Table 3.
Top 10 keywords extracted from LDA topics for Social Science.

Discipline & Topic	Top Keywords	Interpretation
Social Sciences - Topic 0	Approach, model, data, forecasting, image, networks, system, novel, cancer, modeling	Forecasting and social modeling
Social Sciences - Topic 1	Image, network, analysis, model, automatic, networks, design, neural, MIMO, massive	Image and neural network-based analysis
Social Sciences - Topic 2	Approach, model, image, ensemble, modeling, prediction, predicting, method, framework, hybrid	Ensemble/hybrid prediction approaches
Social Sciences - Topic 3	Model, forecasting, analysis, framework, data, method, methods, approach, power, algorithm	Algorithmic models for policy/power data

Table 3 summarizes the 10 main keywords for each topic, and Figure 3 in the appendix visualizes the results. While the coherence score (0.4709) was lower than that of the other two domains, this result should be interpreted in light of the epistemological heterogeneity characteristic of social sciences. Unlike engineering and natural sciences, which often deal with well-structured and quantitatively robust data, social sciences engage with more ambiguous, interpretive, and contextual information. Thus, lower coherence may reflect thematic diversity rather than model weakness. Keywords such as *framework*, *forecasting*, and *modeling* underscore a growing methodological convergence toward data-centric inquiry [17, 18].

Importantly, the appearance of the term *algorithm* across multiple topics reflects a deeper methodological transformation. Social scientists are increasingly shifting from traditional qualitative techniques such as ethnography, interviews, and archival analysis toward approaches that mine, classify, and predict behavior using unstructured digital traces. These traces include user-generated content from social media, GPS and mobility data from smartphones, and interaction patterns from online platforms. The capacity of deep learning to analyze such data in real time creates new opportunities for evidence-based decision-making, policy intervention, and behavioral simulation.

This represents a paradigm shift in social science research, where AI tools are not only used to analyze existing social systems but also to model hypothetical policy outcomes, explore causal relationships, and predict future societal trends. Moreover, concerns about algorithmic bias, data privacy, and explainability are spurring meta-research into the ethics and governance of AI within the social realm.

In conclusion, the social sciences are undergoing a transitional phase with respect to deep learning adoption. Although still in its formative stages compared to other disciplines, the field is showing promising signs of embracing algorithmic thinking, hybrid modeling, and large-scale data analytics. This evolution signals a move toward computational reflexivity, where the tools of analysis themselves become part of the subject of study. As deep learning matures within social research, it holds the potential to transform how we understand, model, and shape human behavior, institutions, and societal change.

5. Discussion and Conclusions

This study set out to examine the evolving landscape of deep learning research through a cross-disciplinary lens, employing Latent Dirichlet Allocation (LDA) topic modeling to analyze 3,000 research paper titles drawn from engineering, natural sciences, and social sciences. Through this methodologically rigorous and data-driven approach, we identified four latent thematic clusters per discipline, uncovering distinctive research priorities and methodological tendencies in each field. Our results highlight the divergent yet occasionally overlapping ways in which deep learning has been conceptualized, implemented, and adapted across disciplinary boundaries.

In the domain of engineering, the results strongly indicate a focus on pragmatic applications of deep learning. The prominence of topics related to real-time control, embedded neural architectures, and autonomous systems reflects the field's alignment with hardware integration, reliability, and performance optimization. Such trends are emblematic of engineering's emphasis on system-level implementation, where deep learning is increasingly serving as the backbone for automation in areas like robotics, industrial diagnostics, and communication systems. The co-occurrence of terms such as "channel," "monitoring," and "vehicle" underscores this applied orientation.

By contrast, the natural sciences exhibit a pattern centered on empirical precision, clinical integration, and diagnostic automation. Topics such as MRI-based image analysis, dose optimization, and multimodal data assessment indicate the growing maturity of deep learning within biomedical and geophysical contexts. The high coherence scores in this domain support the idea that natural sciences are leveraging AI not merely as an analytical supplement but as a fundamental methodological core. In fields such as radiology, oncology, and seismology, deep learning models are increasingly embedded in scientific workflows that demand interpretability, reproducibility, and high-dimensional data handling.

The social sciences, in comparison, present a broader and more pluralistic thematic landscape. The diversity of topics from ensemble forecasting models to algorithmic governance reflects an ongoing shift toward computational social science. While the coherence scores are comparatively lower, this aligns with the field's interpretive flexibility and epistemological heterogeneity. Importantly, terms such as "policy," "behavior," and "algorithm" suggest a growing engagement with the ethical, societal, and human-centered implications of artificial intelligence. In this context, deep learning is not simply a predictive instrument but a lens through which social systems, norms, and inequalities are being re-examined.

Notwithstanding these disciplinary distinctions, several cross-cutting patterns emerged. The frequent appearance of terms such as "model," "framework," "analysis," and "system" across all domains signals a convergence toward shared methodological practices. Deep learning appears to be catalyzing an epistemic realignment in which modeling and data-driven reasoning act as unifying paradigms. This not only facilitates interdisciplinary dialogue but also enables methodological transfers such as applying diagnostic tools from medical imaging to behavioral prediction in policy settings. Such convergence underscores the growing importance of computational thinking as a lingua franca across sciences.

From a methodological perspective, the use of LDA topic modeling proved to be an effective meta-analytical strategy for synthesizing scientific discourse. Beyond identifying thematic prevalence, this approach illuminated the internal organization of research domains and revealed how terminological clusters evolve and cohere. Evaluation metrics like coherence and perplexity provided quantitative measures of topic quality, while qualitative interpretation allowed for context-sensitive domain mapping. This dual-mode approach reinforces the value of unsupervised learning in bibliometric and science-of-science studies [19, 20].

Looking ahead, future research may benefit from several extensions. First, incorporating full-text analysis rather than title-only data would enrich the semantic granularity of topic modeling, allowing for deeper insights into argumentation patterns and methodological nuances. Second, moving beyond disciplinary groupings to explore subdomain-level differences (e.g., within engineering: energy systems vs. robotics) could yield a finer-grained understanding of research specialization. Third, integrating LDA with citation network analysis, co-authorship data, or author-topic modeling would illuminate patterns of collaboration, intellectual lineage, and knowledge diffusion. Finally, the inclusion of human-in-the-loop feedback mechanisms could help refine topic validity and enhance model interpretability for interdisciplinary audiences.

In summary, this study affirms that deep learning is not merely a technological phenomenon but a transformational force that is reshaping scientific inquiry across domains. It acts as both an enabler of discovery and a catalyst for methodological innovation. By adopting a comparative framework, we offer a novel perspective on how different academic communities are internalizing and operationalizing AI technologies. Ultimately, this work contributes to a broader conversation on the epistemological and institutional implications of machine learning, encouraging further meta-analytical research that bridges disciplinary divides and fosters integrative knowledge production.

Funding:
This paper was supported by Joongbu University Research & Development Fund, in 2025.

Transparency:
The author confirms that the manuscript is an honest, accurate, and transparent account of the study; that no vital features of the study have been omitted; and that any discrepancies from the study as planned have been explained. This study followed all ethical practices during writing.

Acknowledgments:
During the preparation of this manuscript/study, the author(s) used ChatGPT 4o for the purposes of Proofreading sentences.

References
[1] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent Dirichlet allocation," *Journal of Machine Learning Research*, vol. 3, pp. 993–1022, 2003.
[2] T. L. Griffiths and M. Steyvers, "Finding scientific topics," *Proceedings of the National Academy of Sciences*, vol. 101, no. suppl_1, pp. 5228-5235, 2004. https://doi.org/10.1073/pnas.0307752101
[3] A. McCallum, "MALLET: A machine learning for language toolkit," 2002. http://mallet.cs.umass.edu
[4] R. Rehurek and P. Sojka, "Gensim—statistical semantics in Python," 2010. https://radimrehurek.com/gensim/
[5] C. Wang and D. M. Blei, "Collaborative topic modeling for recommending scientific articles," in *Proceedings of the 17th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '11) (pp. 448–456). New York, NY, USA: ACM*, 2011.
[6] C. Sievert and K. Shirley, "LDAvis: A method for visualizing and interpreting topics," in *Proceedings of the Workshop on Interactive Language Learning, Visualization, and Interfaces*, 2014.
[7] L. Yu, Y. Wu, Y. Wang, and J. Zeng, "A hybrid approach for academic topic detection in research papers," *Expert Systems with Applications*, vol. 127, pp. 245–258, 2019.
[8] J. Chuang, D. Ramage, C. D. Manning, and J. Heer, "Interpretation and trust: Designing model-driven visualizations for text analysis," in *Proceedings of the 2012 CHI Conference on Human Factors in Computing Systems (pp. 443–452). New York, NY, USA: ACM*, 2012.
[9] C. Jacobi, W. van Atteveldt, and K. Welbers, "Quantitative analysis of large amounts of journalistic texts using topic modelling," *Digital Journalism*, vol. 4, no. 1, pp. 89-106, 2016. https://doi.org/10.1080/21670811.2015.1093271
[10] S. Bird, E. Klein, and E. Loper, *Natural language processing with Python.* Sebastopol, CA: O'Reilly Media, 2009.
[11] S. Syed and M. Spruit, "Full-text or abstract? examining topic coherence scores using latent dirichlet allocation," in *2017 IEEE International Conference on Data Science and Advanced Analytics (DSAA) (pp. 165-174). IEEE,* 2017.
[12] J. Tang, Z. Meng, X. Nguyen, Q. Mei, and M. Zhang, "Understanding the limiting factors of topic modeling via posterior contraction analysis," in *Proceedings of the 31st International Conference on Machine Learning (pp. 190–198). PMLR*, 2014.

[13] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean, "Distributed representations of words and phrases and their compositionality," in *Proceedings of the 26th Annual Conference on Neural Information Processing Systems (NIPS 2013). Curran Associates, Inc*, 2013.
[14] J. Devlin, M.-W. Chang, K. Lee, and K. Toutanova, "Bert: Pre-training of deep bidirectional transformers for language understanding," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 2019.
[15] J. Schmidhuber, "Deep learning in neural networks: An overview," *Neural Networks*, vol. 61, pp. 85-117, 2015. https://doi.org/10.1016/j.neunet.2014.09.003
[16] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *Nature*, vol. 521, no. 7553, pp. 436-444, 2015. https://doi.org/10.1038/nature14539
[17] R. Alghamdi and K. Alfalqi, "A survey of topic modeling in text mining," *International Journal of Advanced Computer Science and Applications*, vol. 6, no. 1, pp. 147–153, 2015. https://doi.org/10.14569/IJACSA.2015.060121
[18] W. X. Zhao, J. Jiang, J. He, and Y. Song, "A collaborative approach to topic modeling for short texts," in *Proceedings of the 21st ACM International Conference on Information and Knowledge Management (CIKM '12) (pp. 1123–1132). New York, NY, USA: ACM*, 2012.
[19] I. Goodfellow, Y. Bengio, and A. Courville, *Deep learning*. Cambridge, MA: MIT Press, 2016.
[20] Y. Liu *et al.*, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019. https://doi.org/10.48550/arXiv.1907.11692

Appendix A.
Topic Visualization by Domain.

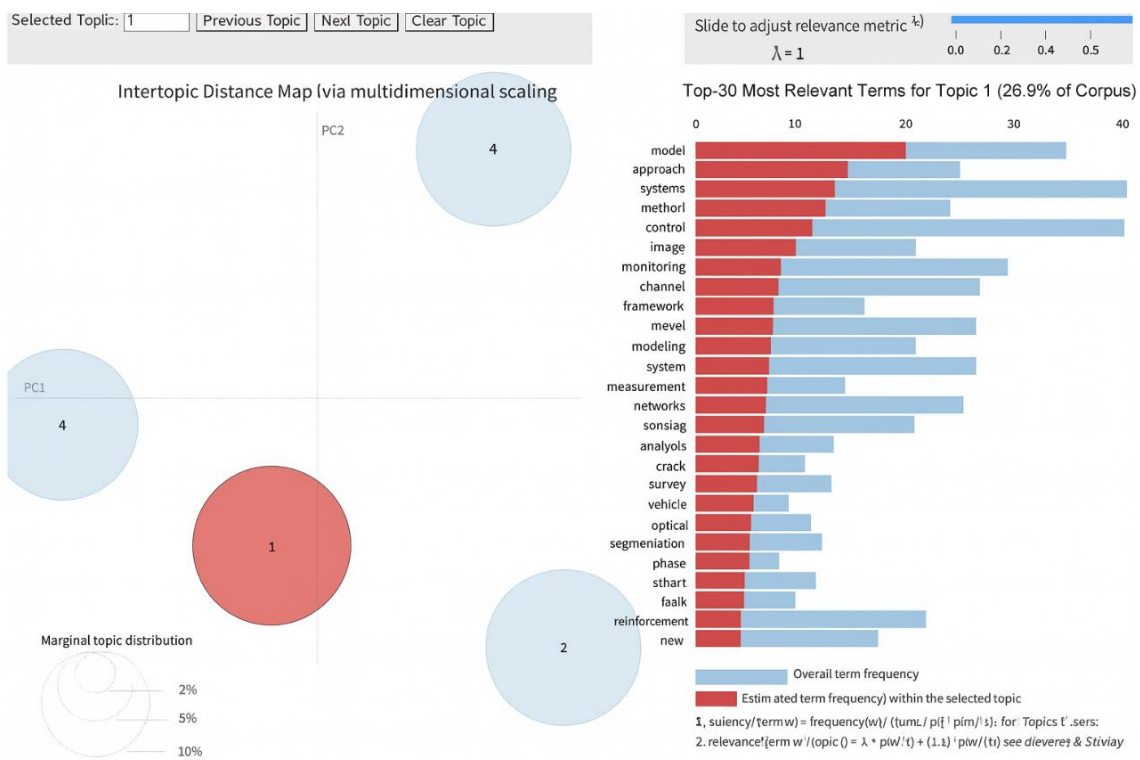

Figure 1.
Intertopic Distance Map and Top Terms for the Engineering Domain.

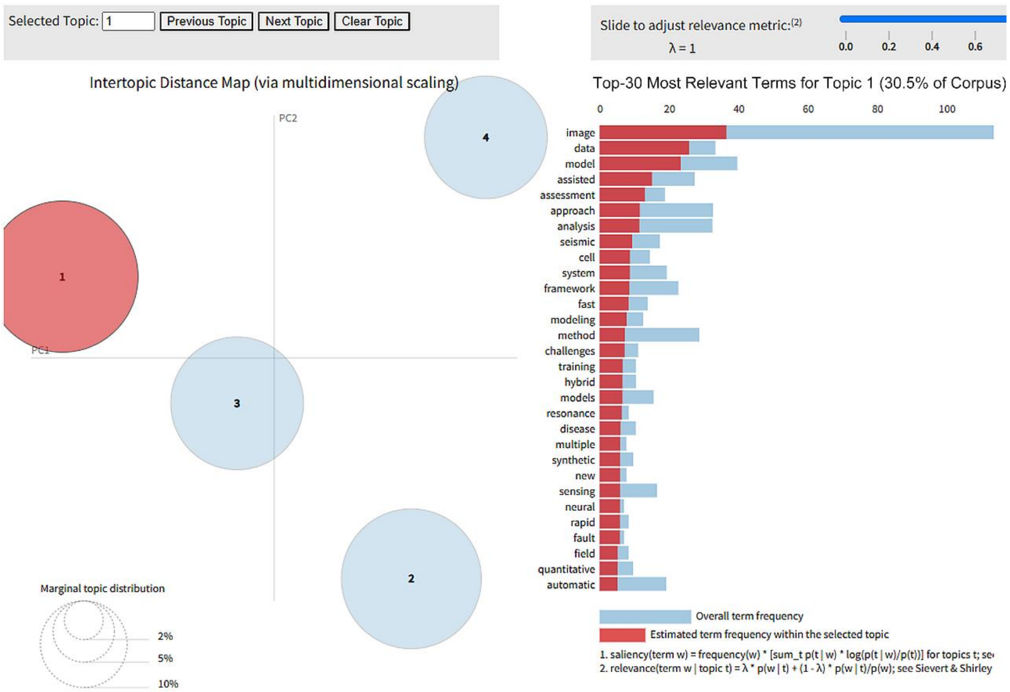

Figure 2.
Intertopic Distance Map and Top Terms for Natural Sciences Domain.

ISSN: 2576-8484
Vol. 9, No. 12: 66-80, 2025
DOI: 10.55214/2576-8484.v9i12.11273

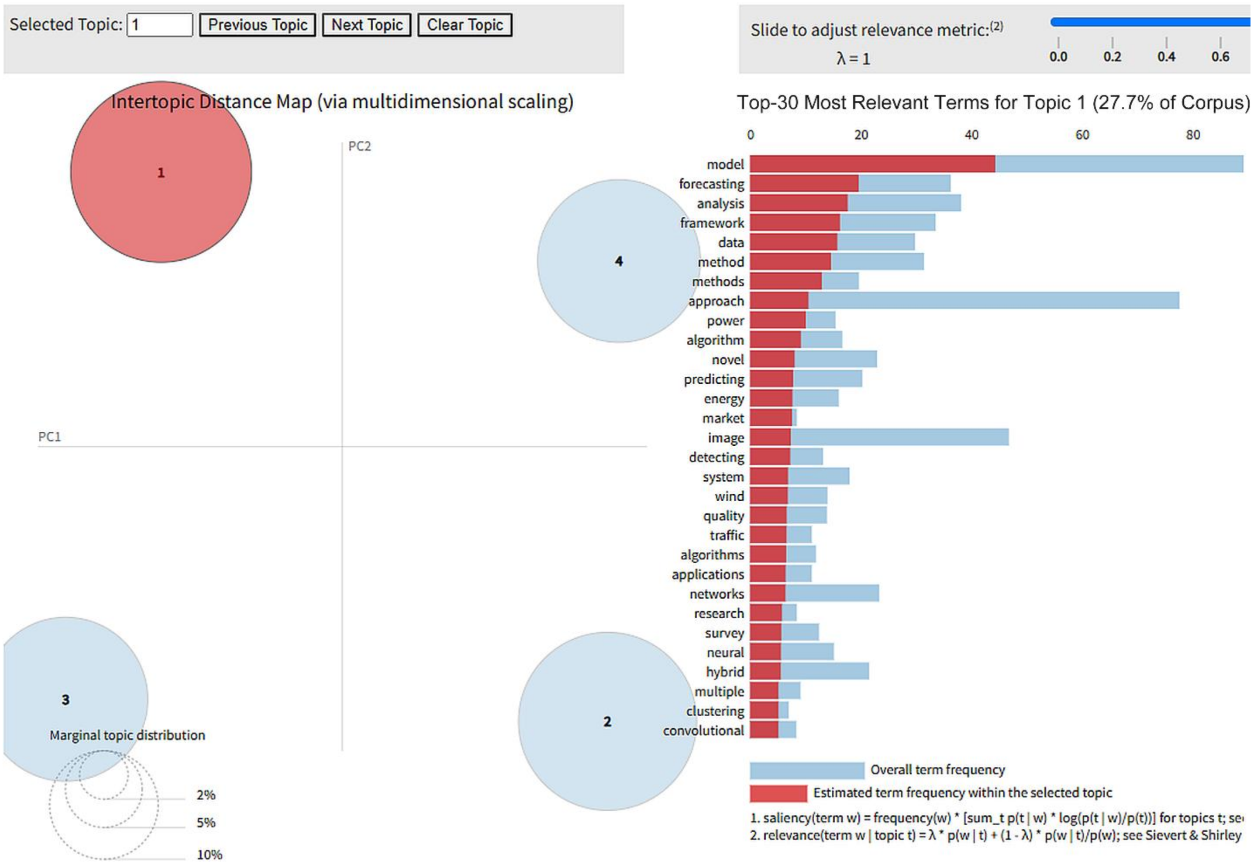

Figure 3.
Intertopic Distance Map and Top Terms for Social Sciences Domain