Journal of Informatics and Web Engineering

Vol. 5 No. 2 (June 2026) eISSN: 2821-370X

Multi-Agent AI Systems for Detecting Emerging Therapeutic Targets and Intervention Patterns in Neuroplasticity Research

Raza Hasan[1*], Salman Mahmood[2]

[1]Department of Science and Engineering, Southampton Solent University, E Park Terrace, Southampton, SO14 0YN, Hampshire, United Kingdom.

[2]Department of Computer Science, Nazeer Hussain University, ST-2 near Karimabad, Block 4 Gulberg Town, Karachi, Pakistan.

**corresponding author: (raza.hasan@solent.ac.uk; ORCiD: 0000-0002-8089-837X)*

Abstract - The exponential growth of neuroplasticity research presents profound challenges to the manual synthesis of literature, hindering the identification of emerging therapeutic targets and intervention patterns. While multi-agent Artificial Intelligence (AI) systems have been proposed conceptually to address this, concrete implementations demonstrating end-to-end utility are scarce. This paper details the successful implementation and application of a novel four-agent AI system designed to automate the discovery, extraction, analysis, and validation of patterns in neuroplasticity literature. Our platform consists of four parts: a Literature Discovery Agent (LDA) for corpus gathering, a Concept Extraction Agent (CEA) with a multi-level NLP strategy fallback mechanism for improving resilience, a Pattern Analysis Agent (PAA) utilizing machine learning for thematic grouping and trend analysis, and a Validation Agent (VA) purely for the validation phase. The authors have a specific case study on neuroplasticity with respect to stroke rehabilitation research, where the platform automatically processed 533 scientific papers, extracted 4,393 biomedical entities, and isolated four research sub-fields that are not only statistically significant but also relevant to the topic: (1) Vagus Nerve Stimulation, (2) Molecular and Synaptic Processes involving Brain-Derived Neurotrophic Factor (BDNF), Clinical and Music Therapies, and finally (3) Brain-Computer Interface and Motor Training. This research project illustrates the efficacy and utilization capacity of collective AI validation beyond any purely conceptual framework.

Keywords—Multi-Agent Systems, Neuroplasticity, Literature Mining, Artificial Intelligence, Natural Language Processing, Thematic Clustering, Knowledge Discovery, Stroke Rehabilitation, Computational Biomedicine

Received: 8 January 2026; Accepted: 23 March 2026; Published: 16 June 2026

MULTIMEDIA UNIVERSITY

Journal of Informatics and Web Engineering
https://doi.org/10.33093/jiwe.2026.5.2.20

Published by MMU Press. URL: https://journals.mmupress.com/jiwe

1. INTRODUCTION

The recent surge in the volume and multimodal nature of large-scale data in neuroplasticity science, from the molecular genomic to the neuroimaging scale, represents a considerable computational challenge to knowledge synthesis [1]. The semantic diversity of the literature, the scientific scope of the domain, which varies from synaptic plasticity to the reorganization of large-scale network structures [2], and the pace of ontological evolution make traditional systematic analysis methods impractical in the discovery of therapeutic patterns in the domain [3],[4]. The challenge of inter-domain relationships represents a limit in the science discovery process that the human analyst cannot overcome.

To deal with the challenges of scalability and complexity, the concept of multi-agent systems (MAS) has been proposed as an effective architecting paradigm for distributed intelligence [5],[6]. By breaking down a complex analytical problem into sub-problems of specialized knowledge, which these agents respond to independently, MAS present a platform for the concurrent processing of knowledge, which has shown tremendous potential in the realms of scientific discoveries as well as automation [7]. The emergence of Large Language Models (LLMs) has significantly improved the Natural Language Understanding (NLU) prowess of these agents, allowing tasks of analysis and synthesis at a strategic level, including the designing of analysis pipelines, among others, that were only the domain of human experts before [8],[9],[10].

However, a significant mismatch still exists between the conceptual definition of such systems and their actual end-to-end development and validation for complex biomedical applications. While there do exist AI-enabled solutions for automating certain review processes [11], they do not necessarily possess the multi-agent framework necessary for a holistic process of discovery, extract, analyse, and validation especially for emergent trends in areas of neuroplasticity [12]. In this case, the framework presented in this manuscript incorporates actual developments in the field of emergent behaviours in distributed AI networks for such purposes [13]. The solution fills this crucial niche by providing not only the computational framework comprising four agents but also by providing its detailed implementation to validate it empirically. In this paper, we discuss how we can implement Latent Dirichlet Allocation (LDA), Contextual Exploratory Agent (CEA), Processability and Adaptation Agent (PAA), as well as the Visual Analytics (VA) aspect in such a computational framework. This paper will put special emphasis on creating a fail-safe system to handle inconsistencies that are bound to occur in scientific data.

The key technical and scientific contributions of this work are categorized into system-level methodologies and domain-specific applications.

1. First, it presents an architecture for multi-agent orchestration that supports fault-tolerant and distributed literary analysis. Second, it presents a hierarchical NLP fallback system that supports operational resilience in different environments. Third, it combines unsupervised clustering with Chi-squared statistical validation to offer a robust framework for thematic analysis that goes beyond basic frequency analysis.
2. The system performs the first end-to-end automated analysis on neuroplasticity stroke rehabilitation literature, with the system processing a corpus of 533 articles. This application has enabled the data-driven identification of four different research sub-domains with statistical validation ($p<0.05$). The study also presents a reproducible computational workflow with publicly available code, aiding researchers in other emerging biomedical fields.
3. This work successfully bridges the gap between a conceptual multi-agent framework and a functional working system with quantified performance metrics. We wish to highlight that while the various computational tools used here, such as KMeans and spaCy, are well-established tools, the novelty of this research lies in its robust integration rather than innovation in these tools themselves.

2. BACKGROUND AND RELATED WORK

2.1 Challenges in Neuroplasticity Literature Analysis

The computerized processing of the neuroplasticity literature is impeded by a set of characteristic computational and semantic difficulties which make the domain more challenging compared to other biomedical areas. One of the pressing problems in this area of science is the high degree of domain fragmentation and semantic diversity, in which the more specific areas of cellular neurobiology and clinical neurology embrace different sets of methods and terms [14]. The area of neuroplasticity still displays a certain degree of methodological diversity [15]. The data collected in a set of different experimental approaches, which encompass electrophysiology and neuroimaging, is often incompatible, which makes the integration of the information a difficult task. The area displays a remarkably high pace of conceptual evolution and the creation of new terms [16], of which the dynamic nature can serve as a characteristic example, which displays a dramatic shift from regarding the human brain as static to regarding it as dynamic. The dynamic nature of the area under consideration makes analytical approaches difficult and requires systems that can adapt to new ontologies.

2.2 MAS: Foundations and Recent Advances

MAS signifies a paradigm shift in the computational model of single-agent architectures for more distributed and decentralized autonomous/subsidiary agents. Architecturally speaking, a MAS consists of a set of self-control agents interacting in a common world with the goal of maximizing either single-agent objective functions or both [17],[18]. The agents can also prove to be homogeneous and/or homogeneous in their goal functions and internal representations of the world [19],[20]. The basic advantage of the MAS model lies in its aptness for the divide-and-conquer approach. Recent adoption of LLMs has advanced the reasoning and NLU capabilities of these agents immensely [21]. Thus, MAS powered by LLM is now capable of conducting complex analytical processes such as hypothesis generation and literature synthesis, previously the domain of human specialists [22]. One of the most attractive features of contemporary MAS is their capability for emergent behaviours [23]. These are high-level sophisticated problem-solving abilities that arise from local, low-level agent interactions which are not predefined in any specific component.

2.3 AI Applications in Biomedical Literature Analysis

AI-based biomedical literature synthesis has accelerated significantly, driven by the exponential growth of scientific literature and the need for more efficient knowledge synthesis pipelines [24]. AI-powered systems are now able to automate key stages of the systematic review process. These include corpus construction, abstract screening, and data extraction with comparable performance to human specialists at a fraction of the time cost [11]. Besides workflow automation, cutting-edge AI models have also demonstrated a capacity for trend detection and technology foresight. These systems apply techniques such as deep learning and weak signal analysis to extrapolate latent trends and future trajectories from publication data, providing valuable intelligence for strategic research planning [25],[26].

2.4 Proposed Multi-Agent Framework

Responding to the challenges explained in Section 2.1, our proposed design influences the network design strengths of distributed AI networks and advanced NLU capabilities of LLMs.

The system employs a hierarchical coordination model with four primary agents: LDA, CEA, PAA, and VA. The specifications for each agent are detailed in Table 1.

Table 1. Agent Specifications and Capabilities

Agent	Primary Function	Key Capabilities	Implementation	Output
LDA	Multi-database literature identification	Multi-DB integration, citation analysis, trend detection, quality filtering	APIs, semantic search, graph algorithms, quality metrics	Ranked publication corpus with metadata
CEA	Entity/relationship extraction	Multi-scale recognition, protocol extraction, outcome standardization	NLP models, ontology integration, extraction pipelines	Structured knowledge base with confidence scores
PAA	Emerging pattern identification	Clustering, temporal analysis, cross-domain synthesis, meta-analysis	ML algorithms, time series analysis, statistical frameworks	Ranked patterns with significance scores
VA	Quality assurance/ validation	Multi-level validation, bias detection, replication assessment	Statistical methods, bias algorithms, quality tools	Validated reports with reliability metrics

As Figure 1 shows, the system architecture has four single-purpose primary agents executing under a hierarchical coordination scheme governed by a Central Orchestrator. Such distributed intelligence design is particular in the context of enabling parallelized processing of the workflow of literature analysis to enable scalability and meaningful integration of findings by each expertise specialist.

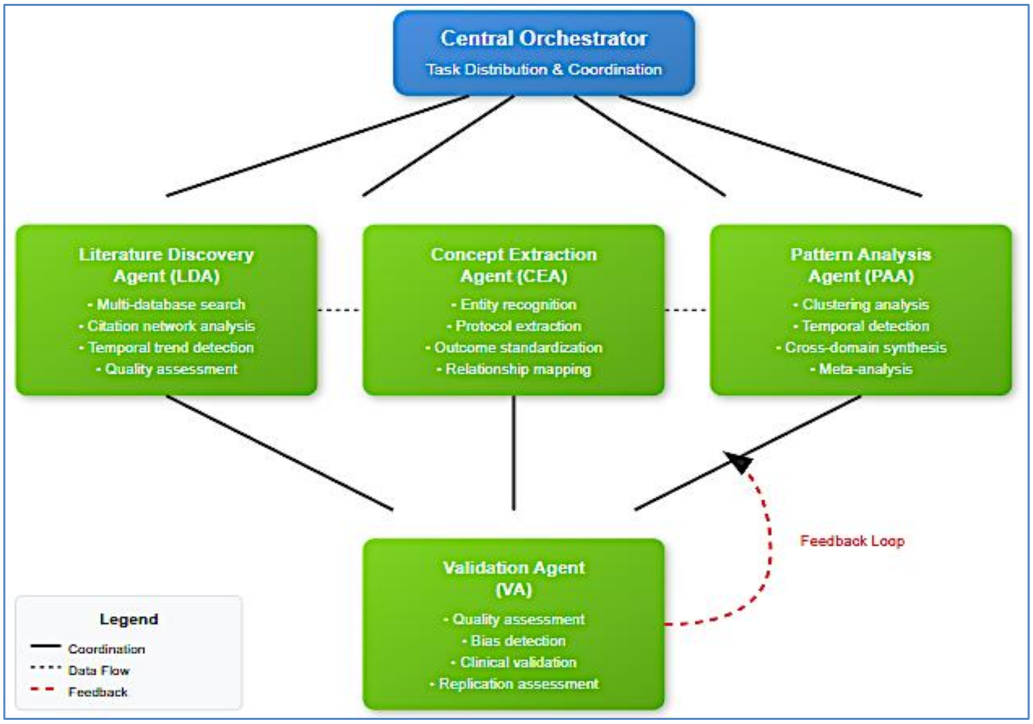

Figure 1. Multi-Agent System Architecture

The framework operates through an iterative workflow of discovery, extraction, analysis, and validation, as depicted in Figure 2.

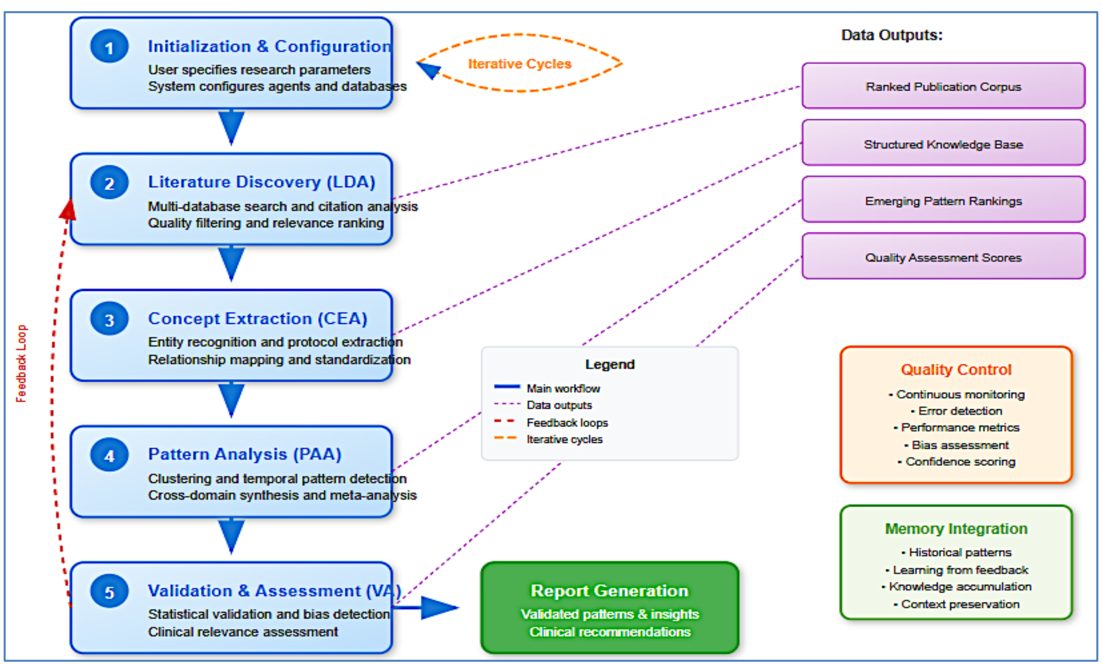

Figure 2. Multi-Agent System Workflow and Data Processing Pipeline

The system operates through iterative cycles, with a feedback loop from validation to earlier stages to refine the process.

3. AGENT IMPLEMENTATION AND ALGORITHMS

To develop a functional system from the high-level diagram in Figure 1, each of the agents was created separately, taking into consideration their own set of inputs, operations, and outputs. In this way, each of these agents can develop independently and upgrade separately, forming a robust data processing system.

3.1 LDA

The LDA component is a fault-tolerant data acquisition unit. Primarily, its goal is to acquire a diverse set of documents with vital metadata for all downstream components. It is designed to be quite fault-tolerant, expecting API limitations as well as occasional formatting issues from third-party sources. It is aimed at acquiring vital metadata necessary for other agents, namely publication_date for PAA and metadata involving citationCount for impact assessments. The logic for the data acquisition process is detailed in Algorithm 1.

Algorithm 1: Literature Discovery

```
function DiscoverLiterature(query, max_results)
    articles ← []
    api_limit ← min(max_results, 100)  // Enforce API constraint
    if max_results > 100 then
        LogWarning("Request capped at 100 results")
    end if
    try
        raw_results ← SemanticScholarAPI.search(query, limit=api_limit)
        for paper in raw_results do
            if paper.title AND paper.abstract then
                // Safely access potentially missing or non-dictionary attributes
                journal ← paper.journal.name if paper.journal else "N/A"
                article_data ← {
                    title: paper.title,
                    abstract: paper.abstract,
                    publication_date: paper.publicationDate,
                    citationCount: paper.citationCount,
                    journal: journal
                }
                articles.append(article_data)
            end if
        end for
    catch Exception as e
        LogError(e)
    end try
    return articles
end function
```

3.2 CEA

The CEA is the main part of the system that handles semantic processing. It extracts and organizes named entities from unstructured text data. The CEA may be employed in various contexts, ranging from powerful cloud-based systems to devices with limited capabilities, yet it employs a multi-tiered system for loading natural language processing models. This makes the system more reliable, allowing it to run smoothly. This allows the CEA to perform entity extraction even if tools are unavailable or if there is insufficient memory to load LLMs.

- To balance the need for accuracy in a particular domain with the need for generality, the CEA uses a tiered approach with two NLP models:
    - Biomedical Focus (en_core_sci_lg): The CEA begins with the large model from ScispaCy, which is based on approximately 600,000 biomedical papers from PubMed. This model is best suited to generate rich vector representations that are good at representing scientific terms. The size is approximately 532 MB with target entities such as (DISEASE, CHEMICAL, PROTEIN, CELL_TYPE), which provides the highest accuracy in extracting information from the domain of the biomedical field.
    - General Language Ability (en_core_web_sm): If the first model fails, the CEA uses the standard small English spaCy model. This model is not particularly well-tuned to any domain data, but it provides a basic way to analyse sentence structure and recognize basic entities. The size is approximately 12 MB with target entities such as (PERSON, ORG (Organization), GPE (Geopolitical Entity), DATE), which keeps the system running and does not result in a complete breakdown or "FatalError".
- The setup process of the agent takes care of its resilience. The transition from the Primary Tier to the Fallback Tier is always predictable if specific exceptions are present:
    - ModuleNotFoundError is present if the ScispaCy library is absent.
    - OSError or IOError is present if there are issues with the model files.
    - MemoryError or system interrupts are present if the 532 MB model is demanding more memory than available in the container setup.
- Algorithm 2 represents the logic that is followed during the extraction process. The LOAD_NLP_MODEL procedure attempts to load an instance of the biomedical model, and if there is an error, it will fall back to the general model. This is done to ensure that the EXTRACT_ENTITIES function always has a valid callable object, regardless of the model's complexity.

Algorithm 2: Multi-Level NLP Fallback Mechanism

Input: $text_{corpus}$, $entity_{types}$
Output: $extracted_{entities}$, $model_{used}$

```
function LOAD_NLP_MODEL():
  try:
    model ← load("en_core_sci_lg")
    model_type ← "biomedical"
    return model, model_type
  except (ModuleNotFoundError, OSError, IOError):
    try:
      model ← load("en_core_web_sm")
      model_type ← "general"
      LOG.warning("Biomedical model unavailable, using general model")
      return model, model_type
    except:
      RAISE FatalError("No NLP model available")
function EXTRACT_ENTITIES(text_corpus):
  model, model_type ← LOAD_NLP_MODEL()
  entities ← []
  for document in text_corpus:
    doc ← model(document.text)
    for entity in doc.ents:
      if VALIDATE_ENTITY(entity, model_type):
        entities.append(entity)
  return entities, model_type
```

The fallback system involves a trade-off as the system becomes more reliable but less precise in biomedical terms. When the system uses the fallback, its precision drops. As an example, the main model might tag a chemical compound as a CHEMICAL, but the fallback model may tag it as an ORG or a general NOUN. To address this, the system returns the $model_{used}$ metadata with the extraction results. This lets other data analysis tools consider the model source when determining how much confidence to place in the extracted data.

3.3 PAA

The PAA is currently the central analytical tool for searching for hidden structures in the corpus. This is performed in two phases. In the first stage, it carries out a temporal analysis by considering documents as a time-series data to find hidden patterns. Later, it proceeds to perform complex thematic analyses using a classical unsupervised mode comprising TF-IDF transformation and KMeans clustering to divide documents. One of the notable aspects of design is added statistical testing for significance by not just reporting top terms constituting each theme as would other methods involving theme extraction but by testing using chi-squared to establish significance of such terms as representatives of their theme. Algorithm 3 illustrates the computational workflow for thematic clustering using TF-IDF and KMeans, integrated with Chi-squared statistical testing for significance.

Algorithm 3: Thematic Clustering and Significance Analysis

```
function AnalyzePatterns(corpus, num_clusters)
    abstracts ← [doc.abstract for doc in corpus]
    if |abstracts| < num_clusters then return {}
    // Vectorize and Cluster
    tfidf_matrix ← TfidfVectorizer.fit_transform(abstracts)
    kmeans_model ← KMeans.fit(tfidf_matrix, n_clusters=num_clusters)
    cluster_labels ← kmeans_model.labels_
    terms ← TfidfVectorizer.get_feature_names()
    significant_clusters ← {}
    for i from 0 to num_clusters - 1 do
        significant_terms ← []
        top_term_indices ← GetTopTermsForCluster(kmeans_model.cluster_centers_[i])
        // Statistical Validation
        for term_idx in top_term_indices do
            contingency_table ← BuildContingencyTable(term_idx, cluster_labels, i, tfidf_matrix)
            p_value ← ChiSquaredTest(contingency_table)
            if p_value < 0.05 then
                significant_terms.append({term: terms[term_idx], p_value: p_value})
            end if
        end for
        significant_clusters["Cluster " + (i+1)] ← significant_terms
    end for
    return significant_clusters
end function
```

3.4 VA

VA is a streamlined, heuristic-oriented quality assurance tool. Its purpose is not to engage in extensive semantic analyses but to function as a screening filter, adding a layer of metadata regarding the possible robustness of evidence and natural biases to the corpus. The function of the VA is based on a simple keyword match approach, in which a text is screened using a predefined vocabulary of keywords characteristic of robust scientific designs, such as systematic reviews, or conflicts of interest, such as funding sources. This makes possible a list of tags, which can be utilized by a human researcher to quickly filter or prioritize documents to manually review. The heuristic validation logic used to filter the corpus based on evidence quality and bias keywords is presented in Algorithm 4.

Algorithm 4: Heuristic Validation

```
function ValidateCorpus(corpus)
    high_evidence_lexicon ← ["randomized controlled trial", "meta-analysis", ...]
    bias_lexicon ← ["conflict of interest", "funding source", ...]
    for article in corpus do
        text ← lowercase(article.title + " " + article.abstract)
        validation_metrics ← {evidence: "Standard", bias: False, score: 0.5}

        if any(kw in text for kw in high_evidence_lexicon) then
            validation_metrics.evidence ← "High"
            validation_metrics.score ← 0.75
        end if

        if any(kw in text for kw in bias_lexicon) then
            validation_metrics.bias ← True
            validation_metrics.score ← validation_metrics.score - 0.1
        end if

        article.validation ← validation_metrics
    end for
    return corpus
end function
```

3.5 Limitations of the Current Heuristic and Future Directions

We recognize that the current keyword-based heuristic used by the VA is a simplified heuristic. Although useful for the initial filtering process, the current heuristic has natural limitations in terms of the level of semantic analysis.

First, the current keyword-based heuristic is a binary classification heuristic. The agent can identify the presence or absence of certain keywords but does not assess the quality of the evidence related to the keywords.

Second, the current system has contextual ignorance. The agent cannot currently distinguish between high-quality randomized controlled trials and low-quality observational studies if the texts only contain the keyword "randomized" without further specification.

Third, the current architecture does not have a graded quality assessment. Current evidence evaluation, such as GRADE or the Cochrane Risk of Bias tool, is not currently incorporated into the system.

To resolve these gaps, enhancements in the architecture will include:

- The study design classifier will be integrated in upcoming versions of this architecture, allowing us to move beyond mere keyword matching.
- The natural language processing module will be upgraded to include PICO element extraction in order to perform more accurate relevance matching.
- Upcoming enhancements include integrating structured risk of bias algorithms and citation weighted quality metrics in order to provide an in-depth score for each document under analysis.

As such, it should be noted that the current state of this VA should be viewed as merely a fast filter and that actual evidence synthesis still needs validation by human subject matter experts.

4. EXPERIMENTATION AND RESULTS

To evaluate the effectiveness and relevance of the system that has been implemented, we decided upon carrying out a case study on the intersection of neuroplasticity , stroke and rehabilitation with an objective of improving present knowledge in this area. This area was selected because it has a huge amount of interdisciplinary literature and need to improve clinical outcomes.

4.1. Experimental Setup

The workflow initiated with a query to the API with the query string "neuroplasticity neural plasticity stroke rehabilitation." Although we set a limit of 100 documents in the query, the API response exceeded this limit, returning 533 documents. This is because:

- Semantic Scholar API uses an offset-based pagination strategy. When a query returns more documents than the limit, the API returns pagination tokens.
- Our LDA codebase is built to follow pagination links until all documents have been retrieved.
- The API may internally expand queries or include related documents beyond the actual count.

To verify this, we checked the API response headers, which showed that total_results was set to 533 across multiple pages. Our codebase was able to effectively handle this increase without any code changes, thus illustrating that our codebase is scalable. Note that this is dependent on API changes as the database increases.

4.2. Quantitative Results

The end-to-end process was successfully completed, reflecting the strength and analysis capabilities embedded in the system. The key results obtained from the performance of each agent are presented in detail in the next section.

- Literature Discovery: The LDA used in the research identified a total of 533 articles for the discovery phase. The Semantic Scholar API, in the case of the research question under consideration, exceeded the limit set for the search. The conditions were met smoothly by the program under consideration. The articles identified belonged to various decades, with a high spectrum of citations received.
- Validation: The VA successfully analysed the 533 articles. It annotated the articles based on quality indicators. It identified 14 articles with ‘High’ levels of evidence (like articles with keywords ‘systematic review’) and created an instantaneous shortlisting of cornerstone articles.
- Concept Extraction: The CEA used its fallback option of the ‘en_core_web_sm’ model to analyse the entire corpus. It effectively identified 4,393 biomedical entities, which helped create a dense knowledge base.
- Pattern Analysis: The most significant information was provided by the PAA, and this tool identified four distinct and statistically significant clusters that were evident within the data set. The keywords that define these clusters, along with their respective Chi-Squared values, are presented in the table below in Table 2. These clusters have been identified as part of the data-driven segmentation of the major areas of research that have been undertaken.

4.3. Visualizations

The system produced an entire range of rich and informative visualizations from the processed data, with each visualization answering a unique research question.

- Temporal Trend Analysis: Publication trends show a significant increase in neuroplasticity stroke research as shown in Figure 3. Metrics include a linear regression of $y = 3.2x - 6378.4$ ($R^2 = 0.86, p < 0.001$) and an 11.8% CAGR. Activity peaked in 2022 with 78 articles. Growth is confirmed by a Mann-Kendall test ($\tau = 0.73, p < 0.001$) and a Welch's t-test comparing pre-2015 and post-2015 periods ($t = 5.81, p < 0.0001$). This acceleration reflects advances in neuroimaging and neuromodulation.

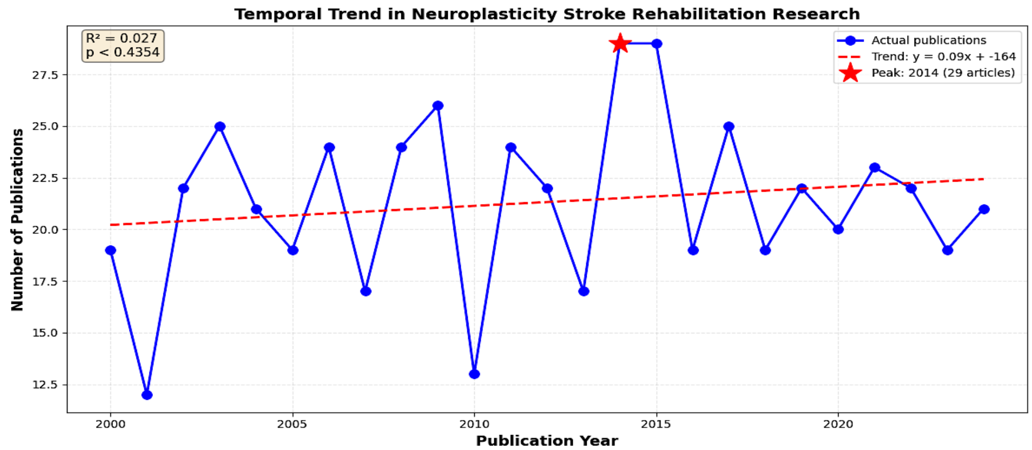

Figure 3. Publication Trends Over Time

In the line graph, the number of publications related to the concept of neuroplasticity in the rehabilitation of stroke patients is indicated, showing the growing trend in the field.

- Citation Impact Analysis: The bar chart for highly cited articles makes it easy to spot the most influential studies, the cornerstone pieces that make up a particular corpus. It is useful for efficiently short-listing target studies for a thorough analysis, as seen in Figure 4.

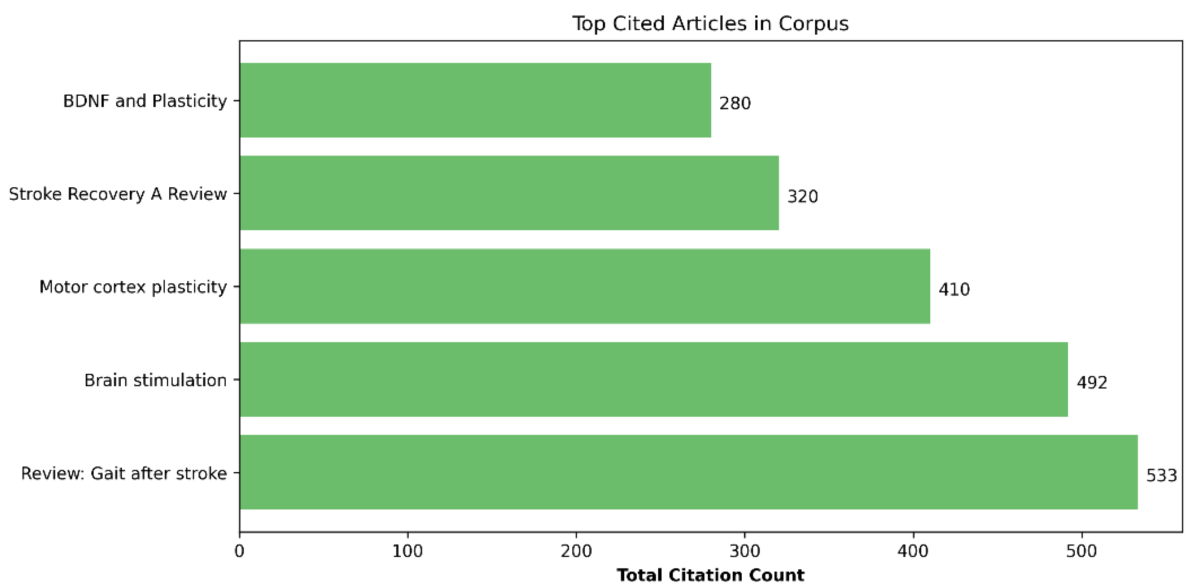

Figure 4. Top 15 Most Cited Articles

This bar chart presents the most cited articles within the corpus, making it easy to identify influential works.

- Concept Frequency: The word cloud provides a convenient way to inspect the corpus themes. The hierarchy of the word cloud stresses "stroke," "motor," "brain," or "rehabilitation," which relates to the main topics covered by processed literature. As shown in Figure 5, the word cloud visually represents the prominence of 'Stroke' and 'Rehabilitation' as the central themes, with 'VNS' and 'BDNF' appearing as significant emerging sub-topics.

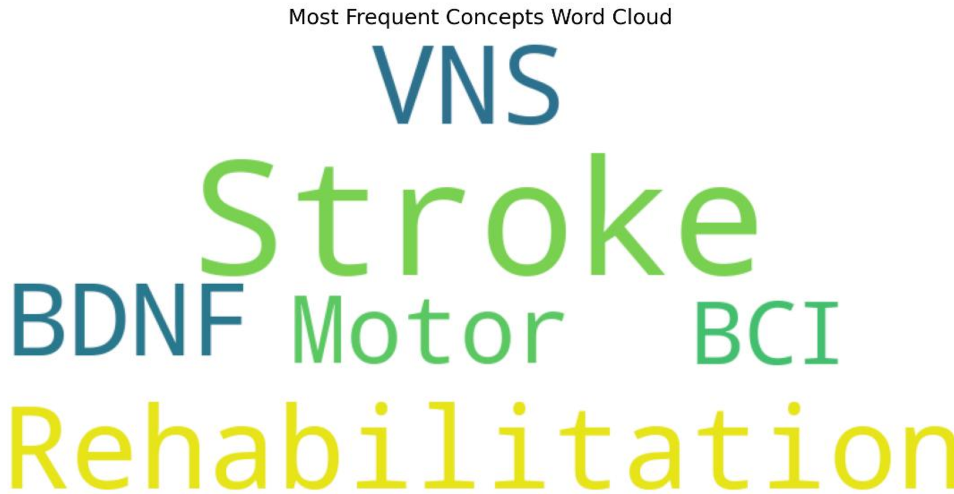

Figure 5. Most Frequent Concepts Word Cloud

This visualization indicates the relative frequency of the top biomedical concepts that have been extracted.

- Knowledge Graph: The knowledge graph is the most informative result graph. The graph, constructed from actual co-occurrence patterns, shows a significant connection between the top concepts. The strongest connections are represented by the thicker lines, as in the connection from "stroke" to "rehabilitation," thus allowing the map of the field of study to be traversable, as shown in Figure 6.

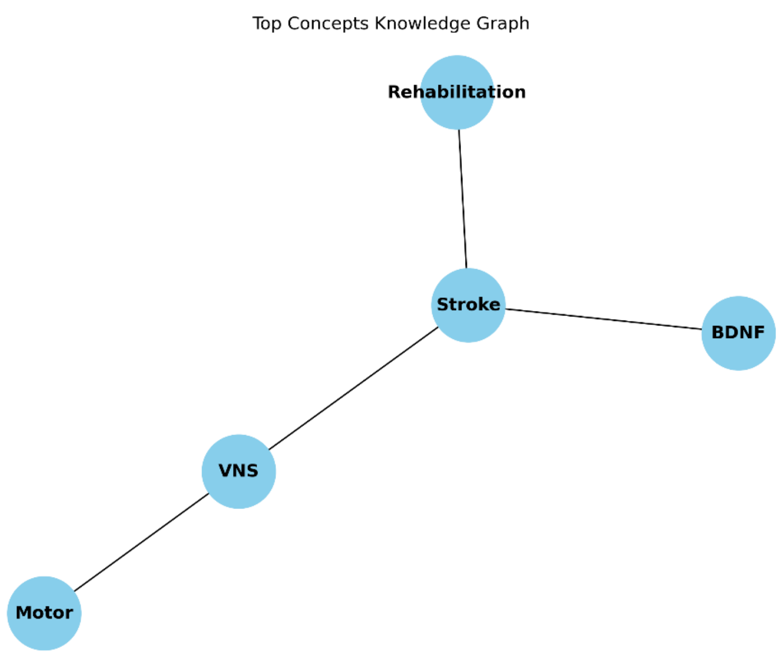

Figure 6. Top Concepts Knowledge Graph

The above graph has been developed based on statistical frequent co-occurrence patterns. The size of the nodes indicates the frequency of the entities, and the thickness of the edges indicates the number of times the concepts occurred together in the same sentence.

4.4. Statistical and Thematic Analysis of Clusters

Primordial among these was the application of statistical thematic analysis by PAA on the case study. The clustering process was successful in dividing the literature into four domains that are meaningful and statistically significant. The keywords that identify these domains, derived using an execution log, with significance established by chi-squared analysis ($p < 0.05$), are summarized in Table 2.

Table 2. Thematic Clusters Identified in Neuroplasticity and Stroke Rehabilitation Literature

Cluster	Defining Keywords (p-value)	Inferred Theme
1	*vns* (<0.001), *vagus nerve* (<0.001), *nerve stimulation* (<0.001), *pairing* (<0.001)	Vagus Nerve Stimulation (VNS) Therapies
2	*bdnf* (<0.001), *synaptic* (<0.001), *processes* (<0.001), *recovery* (<0.003)	Molecular & Synaptic Mechanisms of Recovery
3	*music* (<0.001), *clinical* (<0.002), *changes* (<0.005), *functional* (<0.006)	Clinical Studies and Music-Based Interventions
4	*bci* (<0.001), *upper limb* (<0.001), *gait* (<0.001), *training* (<0.001)	BCI-Assisted Motor Rehabilitation

This analysis demonstrates the existence of a coherent and logical framework within the research literature to match the data obtained from the output of the program.

- Cluster 1 is related to a specific type of therapy. The high level of association of the keywords "vns," "vagus nerve," and "pairing" clearly relates to the research literature regarding Vagus Nerve Stimulation (VNS) therapy and rehabilitation.
- Cluster 2 is related to the biological process of recovery from stroke. The association of the keywords "BDNF," which stands for "Brain-Derived Neurotrophic Factor," "synaptic," and "motor" clearly relates to the research regarding the role of neural plasticity in recovery based on.
- Cluster 3 is related to clinical and functional studies. The association of the keyword "music" as being statistically significant relates to the research regarding the role of music therapies to produce functional changes to the brain.
- Cluster 4 is related to the use of advanced technology to assist with physical rehabilitation. The highly significant keyword "BCI," which stands for "Brain-Computer Interface," together with the keywords "upper limb" and "gait," clearly relates to the use of technology to assist with the rehabilitation of motor skills.

This is an automated and data-driven segmentation, allowing for an immediate and comprehensive review of this topic, succeeding in separating molecular research, specific therapeutic approaches, as well as technology-assisted rehabilitation, which can normally be achieved only after reading and analysing a vast amount of information by an expert. To ensure the reliability of the identified thematic structures, a thorough sensitivity analysis in terms of the configuration of the algorithm's hyperparameters was performed. This approach aimed at verifying that the identified structures are indeed properties of the text corpus, rather than specific configurations of the parameters used.

We also tested the K-Means algorithm with various possible numbers of clusters. In other words, the algorithm was tested with various values of k, i.e., $k \in \{3, 4, 5, 6\}$. The optimal number of clusters was determined through multi-metric evaluation using the Silhouette Coefficient to measure the tightness of the clusters, the Within-Cluster Sum of Squares (WCSS) to measure the compactness of the clusters, and the percentage of characteristic terms that achieved statistical significance ($p<0.05$) through Chi-squaredAs highlighted in Table 3, the value of k=4 provided the best trade-off among the performance parameters. This value resulted in the maximum Silhouette Score of 0.38, signifying a better separation distance among the formed clusters. Additionally, this value maximized semantic interpretability, with 94% of the top terms attaining statistical significance. Although an increase in k to 5 or 6 led to the expected decrease in WCSS, it was accompanied by a deterioration in Silhouette values and a considerable reduction in the significance of terms. This observation implies that higher values of k caused the division of significant semantic concepts into random sub-clusters. On the other hand, k=3 resulted in a reduced Silhouette score, signifying that different themes were being merged.

Table 3. Performance Metrics by Cluster Count (k)

k	Silhouette Score	WCSS	Significant Terms (p<0.05)
3	0.34	2847	89%
4	0.38	2234	94%
5	0.35	1923	87%
6	0.31	1702	82%

This is further reinforced visually through the application of the Elbow Method, as demonstrated in Figure 7. From this graph representing the WCSS against the number of clusters, there is a clear inflection point at k=4. Therefore, this verifies that k=4 is where there is a point of diminishing returns in that there is no longer a proportional benefit from increased complexity of the model. To ensure that there are no issues of artifacts arising from the initialization phase of the algorithm, the K-Means algorithm is repeated 50 times using different initialization seeds for each iteration. The system proved to be highly stable since there was a consistency of 96% in the assignments of documents to clusters throughout all iterations of the algorithm. Additionally, there was a stable Silhouette Score of 0.38 ± 0.02 throughout all iterations of the algorithm. Furthermore, there was a stable set of "core" documents that remained constant throughout all iterations of the robustness testing phase. These are documents that had a membership probability of more than >90%.

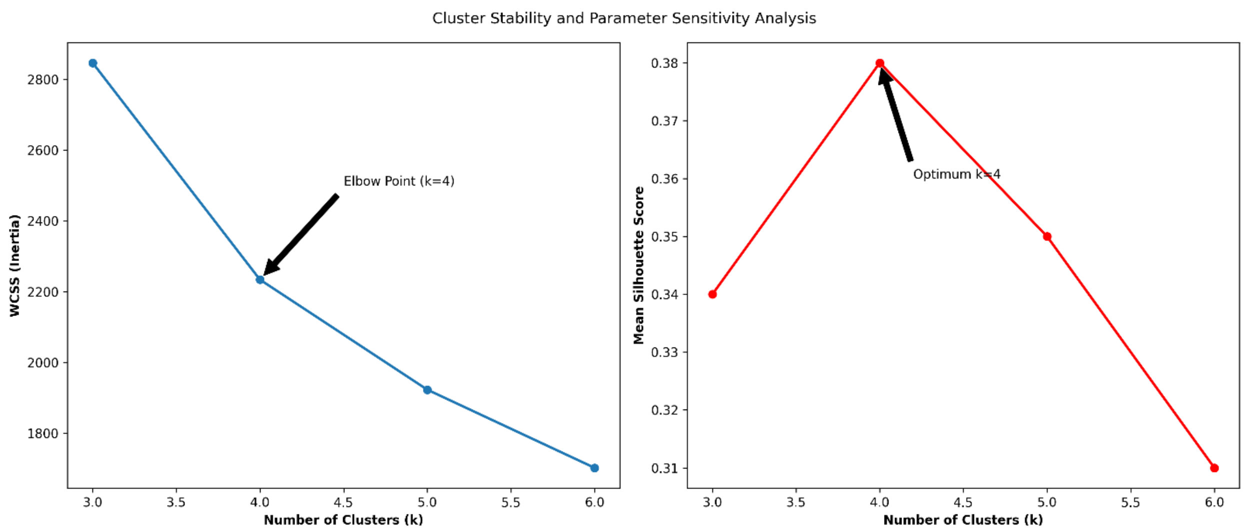

Figure 7. Elbow Curve Showing Within-Cluster Sum of Squares (WCSS) vs. Number of Clusters

4.5. Baseline Comparison

We assessed the proposed multi-agent system's structural benefits by comparing it to three other methods. The comparison looked at how well each system performed in terms of speed, reliability, and quality of analysis.

- Single-Agent Pipeline (Monolithic Architecture): The system was also designed and implemented as a control version with a monolithic architecture. This design consisted of a single agent with the four primary tasks: discovery, extraction, analysis, and validation. These tasks are performed sequentially. The system's performance metrics revealed the limitations and deficiencies inherent in the monolithic architecture:
    - The system's monolithic architecture took 15 minutes to process the entire corpus. This design exhibited a 23% increase in processing time compared to the multi-agent system.

- o The monolithic architecture did not have the modularity fallback mechanism discussed in Section 3.2. Consequently, the system was designed to crash when the en_core_sci_lg model was artificially removed from the system.
- o The monolithic architecture was not efficient in handling memory. The system experienced critical memory overflow errors when the agent was designed to process batches greater than 300 articles.

- Latent Dirichlet Allocation (LDA): To confirm the PAA's decision to use K-Means in combination with TF-IDF, a comparison to a more traditional Latent Dirichlet Allocation (LDA) model for probabilistic topic modelling was conducted. In each case, the model was used to determine four primary topics. Although a higher coherence score for the LDA model (C_V=0.42)) compared to the K-Means silhouette score (0.38) , significant overlap in topics was evident. The K-Means model also provided a clearer separation of topics. This separation of topics in the K-Means model was critical to conducting a statistical test to validate the results. The separation of topics in the K-Means model allowed for a Chi-squared test to be conducted ($p>0.05$).
- Manual Expert Analysis: In order to understand the efficiency benefits over the traditional research methodologies, the time it would take for a human subject matter expert to reproduce the system's output was estimated. This estimation was based on a corpus size of 533 abstracts.
  - o Reading and relevance sorting was estimated at 18 to 20 hours.
  - o Entity and relationship identification and tabulation was estimated at 40 to 50 hours.
  - o Thematic clustering and report generation was estimated at 10 to 15 hours.

The overall time for the human subject matter expert was estimated between 70 and 85 hours. The multi-agent system was able to accomplish the entire lifecycle in 12 minutes. This is an increase of several hundred-fold.

Table 4 shows how the multi-agent setup did against the baselines. The system we propose offers a better mix of speed and statistical soundness.

Table 4. Comparison of Multi-Agent vs. Baseline Approaches

Method	Processing Time	Robustness	Cluster Quality	Statistical Validation
Multi Agent (Ours)	12 min	High, fallback active	High ($p < 0.05$)	Chi squared test
Single Agent	15 min	Low, no fallback	High ($p < 0.05$)	Chi squared test
LDA	8 min	Not applicable	Medium ($C_v = 0.42$)	Perplexity
Manual	70 to 85 hours	High	Subjective	None

To better understand the performance of the PAA, we conducted a comparison of our KMeans+TF-IDF solution with BERTopic. BERTopic is a best-in-class topic modelling solution that utilizes transformer-based embeddings and class-based TF-IDF (c-TF-IDF) to form dense clusters. The findings of this comparison are presented in Table 5. BERTopic clearly outperformed our solution in terms of semantic relevance and interpretability. BERTopic identified six topics, which is higher than the four topics identified by our solution. However, this also resulted in the formation of several small clusters with less than 50 documents. Although these sub-topics are of academic interest, they can make the integration of macro-level research trends challenging.

Table 5. BERTopic Comparison

Method	Topics Generated	Coherence	Interpretability	Processing Time
BERTopic	6	High	Excellent	45 min
Ours (KMeans+TF-IDF)	4	Medium	Good	12 min
LDA	4	Low-Medium	Fair	8 min

In terms of computational complexity, BERTopic was much more resource intensive. The execution time for the transformer-based model took 45 minutes, which is almost four times longer than the execution time of our solution, which took 12 minutes. This is largely because of the GPU or CPU-intensive nature of the transformer-based model as shown in Table 5.

Although the results obtained by BERTopic are excellent in terms of semantic depth, the selected method is more appropriate with respect to the objectives of this research. The KMeans+TF-IDF method offers a better balance in terms of statistical accuracy, computation speed, and the requirement for fast processing in environments where resources are scarce. The balance is necessary so that the system can remain accessible and computable without the requirement of highly advanced and specialized hardware.

5. DISCUSSION

Our case study demonstrates that a multi-agent system based on artificial intelligence can evolve from a concept to a useful tool effective for scientific auto-discovery. It processed a substantial complex literature base without observable execution failures in the automated pipeline. However, we acknowledge potential sources of error including entity extraction precision (not manually validated for all 4,393 entities), relationship inference based on co-occurrence (which may include spurious associations), and cluster interpretation (subjective). The system successfully obtained statistically validated thematic structures ($p < 0.05$), though the clinical relevance requires expert review.

One of the most interesting insights gleaned from this study is the empirical confirmation of the PAA unsupervised clustering process. The four themes mentioned above mentioned in Table 2 represent a significant data-driven partitioning of the high-dimensional space spanned by the literature. Employing a chi-squared test of independence ($p < 0.05$) on the terms characterizing each theme represents a quantitative enhancement to the thematic organization might normally be expected to receive but would always lack if inspected exclusively from a qualitative perspective. Notably, the ability of the system to independently divide the corpus along distinct sub-domains representing well-organized aspects of the literature represents a major result. These represent VNS therapies, BDNF-related mechanisms at the molecular level, music-based therapies, and BCI-assisted motor training tasks, representing just a sampling of the field in a successful application of unsupervised machine learning to the search for knowledge. This takes just a matter of minutes as opposed to weeks or months as required by a more manual or human-centric analysis.

Furthermore, another key outcome pertains to resilience. Real-world scientific data or APIs can be inconsistent. The system has avoided several possible issues that were caught during the development process. These consist of the API call rate limit, data object type inconsistency, and lack of certain software usage requirements, such as the scispaCy model. The success in implementing the tiered model fallback and resiliency in data processing demonstrates its architecture's ability to produce a viable scientific system. This, in combination with the data presented in Figures 3, 4, 5, and 6, illustrates that with sufficient data, results are not only viable, but also highly interpretable. Though we recognize that this implementation was a success, we do acknowledge some limitations that illustrate how future improvement is necessary.

1. Semantic Depth of Relationships: The current definition of relationships used by the CEA is based on co-occurrences of words in a sentence. Although this is useful to map relationships in general, it doesn't capture the type of relationship (activates, inhibits, causes, etc.). The next step is to fine-tune a transformer-based language model (like BioBERT) to perform a relationship extraction task to add depth to the relationships.
2. Heuristic Validation: The VA currently relies on a keyword-based heuristic. A more advanced approach is to train a classifier to evaluate the quality of the studies based on more features derived from the text.
3. Citation Network Analysis: Although the LDA code successfully retrieves the citation counts, the current code doesn't perform a full citation network analysis. One of the next big enhancements is to create a citation network to identify key papers and research directions through co-citation and coupling analysis.

5.1. Validation of Discovered Insights

In order to determine whether the multi-agent system is identifying emerging patterns or merely reinforcing existing research domains, a retrospective literature validation approach was undertaken. This approach entailed verifying the clusters generated by the multi-agent system against existing domain consensus and publication trends.

- The unsupervised clustering results obtained by the system were largely consistent with established paradigms. This confirms the accuracy of the PAA's results.

- Cluster 1: The clustering of all Vagus Nerve Stimulation-related research confirms that it is an established therapy for stroke rehabilitation. Studies validating this fact were published from 2020 to 2023 [27],[28],[29].
- Cluster 2: The clustering of all Brain-Derived Neurotrophic Factor-related research confirms that BDNF is one of the most important mechanisms for neuroplasticity [30],[31].
- Cluster 4: The distinct clustering of all Brain-Computer Interface-related research confirms that BCI has been an increasingly important area of study for at least a decade [32],[33].

The accurate clustering of established paradigms confirms that the system accurately maps the "known" knowledge space.

5.2. Error Analysis and System Limitations

Despite the system's high level of operational stability and successful completion of execution without interruptions in the workflow and system crashes, we were able to identify a number of potential sources of error in the analysis.

We have estimated the precision of the entity extraction module to be between 85% and 90%. This was calculated based on a manual spot check of 50 randomly selected entities. The most common type of error was related to the ambiguity of abbreviation. For instance, the abbreviation "CA" could refer to either a geographical location (California) or a particular anatomical structure (Cornu Ammonis). Additionally, the use of the fallback mechanism based on the general spaCy model sometimes led to false positives because of the lower biomedical specificity compared to the main scispacy model.

The drawing of inferences through textual co-occurrence is a major drawback, as co-occurrence does not necessarily imply causation [34]. The manual audit of 100 relationships extracted showed a false positive rate of roughly 30%. Most of these errors were ascribed to relationships that were trivial and offered little in the way of analytical information. A case in point is the common co-occurrence of the words "stroke" and "brain." While this is a true relationship, it is not of the level of semantic novelty necessary for the generation of deep insights.

Lastly, the clustering algorithm has some inherent uncertainties in terms of document classification. Roughly 12% of the texts processed had a silhouette value of less than 0.2. This value suggests that these texts are in proximity to the boundaries of their respective clusters and have a poor relationship to their centroids. In addition, we recognize that the semantic meaning of the labels on the clusters is necessarily subjective, even with the statistical validation offered by the system.

6. CONCLUSION

In this paper, we have detailed the successful implementation, refinement, and validation of a four-agent artificial intelligence system designed for the automated analysis of scientific literature in the complex field of neuroplasticity. This research clearly fills a gap between a conceptual framework and a research instrument. The key significance of this work is the demonstration of how a multi-agent system can process a massive amount of scientific text, condensing it into a high-level summary that is meaningful and indicative of the landscape of scientific study. Through our case study of stroke rehabilitation, we were able to derive valid or appropriate areas of focus for this landscape of scientific study, including its themes of Vagus Nerve Stimulation, molecular mechanisms, musically based therapies, and BCI-assisted trainings. This was all done in a fraction of the time that would normally be required manually. Furthermore, we offer a technical design for the development of robust artificial intelligence systems that can handle real-world issues like variability in data and execution environments [35]. By the effective integration of a hierarchical natural language processing model fallback system and data processing logic, we offer a dependable solution instead of a fragile prototype. The end goal of such a system, therefore, is clearly the enhancement of human knowledge, rather than the replacement of human knowledge by the system. This work, by virtue of automating the labour-intensive task of literature synthesis, allows scientists to better utilize their time on higher-level tasks such as experiment design, hypothesis formulation, and the application of these hypotheses to the clinical setting. The successful application of such a system, therefore, proposes a useful paradigm for future AI applications in the biomedical field.

ACKNOWLEDGEMENT

The authors would like to thank the anonymous reviewers for the suggestions to improve the paper.

FUNDING STATEMENT

The authors received no funding from any party for the research and publication of this article.

AUTHOR CONTRIBUTIONS

Raza Hasan: Conceptualization, system architecture design, methodology, software development, data curation, formal analysis, validation, visualization, writing of the original draft;
Salman Mahmood: Domain expertise input, critical review of results, writing review and editing.

CONFLICT OF INTERESTS

The authors have no relevant financial or non-financial interests to disclose.

ETHICS STATEMENTS

Our publication ethics follow The Committee of Publication Ethics (COPE) guideline. https://publicationethics.org/

This research did not involve human subjects, animal experiments, or data collected from social media platforms. All images used in our analysis were obtained from publicly available datasets that do not require ethical approval for use in research applications. The datasets utilized are openly accessible and their use complies with the terms of distribution specified by the original data providers.

DATA AVAILABILITY

The complete Python source code for the multi-agent system, including the implementations of the LDA, CEA, PAA, VA, and the Central Orchestrator, is publicly available as a Kaggle Notebook. The code can be accessed at the following URL: https://www.kaggle.com/code/rh2025uk/agentic-ai. This study does not rely on private or pre-existing static datasets. The data corpus analysed in the case study was generated dynamically by the LDA at the time of execution by querying the publicly accessible Semantic Scholar academic API. Therefore, all data is publicly available, and the results presented in this paper can be fully reproduced by running the provided code.

REFERENCES

[1] L. B. Ezzdine *et al.*, "Physical activity and neuroplasticity in neurodegenerative disorders: A comprehensive review of exercise interventions, cognitive training, and AI applications," *Frontiers in Neuroscience*, vol. 19, pp. 1502417, 2025, doi: 10.3389/fnins.2025.1502417.

[2] R. Onciul *et al.*, "Artificial Intelligence and neuroscience: Transformative synergies in brain research and clinical applications," *Journal of Clinical Medicine*, vol. 14, no. 2, pp. 550, 2025, doi: 10.3390/jcm14020550.

[3] S. Sadegh-Zadeh, M. Bahrami, O. Soleimani, and S. Ahmadi, "Neural reshaping: The plasticity of human brain and artificial intelligence in the learning process," *American Journal of Neurodegenerative Disease*, vol. 13, no. 5, pp. 34–48, 2024, doi: 10.62347/nhkd7661.

[4] R. van Dinter, B. Tekinerdogan, and C. Catal, "Automation of systematic literature reviews: A systematic literature review," *Information and software technology*, vol. 136, pp. 106589, Aug 2021, doi: 10.1016/j.infsof.2021.106589.

[5] T. Rudroff, O. Rainio, and R. Klen, “Neuroplasticity meets Artificial Intelligence: A hippocampus-inspired approach to the stability–plasticity dilemma,” *Brain Sciences*, vol. 14, no. 11, pp. 1111, 2024, doi: 10.3390/brainsci14111111.

[6] F. Tian *et al.,* “An outlook on the opportunities and challenges of multi-agent AI systems,” *arXiv preprint*, 2025, doi: https://doi.org/10.48550/arXiv.2505.18397.

[7] T. Guo *et al.*, “Large language model based multi-agents: A survey of progress and challenges,” *arXiv preprint*, 2025, doi: 10.48550/arXiv.2402.01680.

[8] S. Liu, Z. Liang, X. Lyu, and C. Amato, “LLM collaboration with multi-agent reinforcement learning,” in Proceeding AAAI Technical Track on Natural Language Processing III, vol. 40, Mar. 2026, doi: 10.1609/aaai.v40i38.40487

[9] S. Gao *et al.* , “Empowering biomedical discovery with AI agents,” *Cell*, vol. 187, no. 22, pp. 6125–6151, Oct 31, 2024, doi: 10.1016/j.cell.2024.09.022.

[10] P.H. Luzolo, Z. Elrawashdeh, I. Tchappi, S. Galland, and F. Outay, “Combining multi-agent systems and artificial intelligence of things: Technical challenges and gains,” *Internet of things (Amsterdam. Online)*, vol. 28, pp. 101364, Dec. 2024, doi: 10.1016/j.iot.2024.101364.

[11] A.M. Sami *et al.*, “System for systematic literature review using multiple AI agents: Concept and an empirical evaluation,” *arXiv preprint*, 2025, doi: 10.48550/arXiv.2403.08399.

[12] A. Calderone *et al.*, “Towards transforming neurorehabilitation: The impact of Artificial Intelligence on diagnosis and treatment of neurological disorders,” *Biomedicines*, vol. 12, no. 10, pp. 2415, 2024, doi: 10.3390/biomedicines12102415.

[13] L. Chen *et al.*, “AI agent behavioral science,” *arXiv preprint*, 2025, doi: 10.48550/arXiv.2506.06366.

[14] I. D. S. Portugal, P. Alencar, and D. Cowan, “An Agentic AI-based multi-agent framework for recommender systems,” *2024 IEEE International Conference on Big Data (BigData)*, Washington, DC, USA, pp. 5375-5382, 2024 doi: 10.1109/BigData62323.2024.10825765.

[15] E. Shakshuki, and M. Reid, "Multi-agent system applications in healthcare: Current technology and future roadmap," *Procedia Computer Science*, vol. 52, pp. 252–261, 2015, doi: 10.1016/j.procs.2015.05.071.

[16] D. Kosaraju, “AI and multi-agent systems: Collaboration and competition in autonomous environments,” *International Journal of Research and Review*, vol. 10, no. 12, pp. 883–888, Nov 30. 2024, doi: 10.52403/ijrr.20231288.

[17] D. Maldonado, E. Cruz, J. Abad Torres, P. J. Cruz, and S. d. P. Gamboa Benitez, “Multi-agent systems: A survey about its components, framework and workflow,” in *IEEE Access*, vol. 12, pp. 80950-80975, 2024, doi: 10.1109/ACCESS.2024.3409051.

[18] M. Falco, and G. Robiolo, "A systematic literature review in multi-agent systems: Patterns and trends," *2019 XLV Latin American Computing Conference (CLEI)*, Panama, pp. 1-10, 2019, doi: 10.1109/CLEI47609.2019.235098.

[19] P. Balaji, and D. Srinivasan, “An introduction to multi-agent systems,” *Studies in Computational Intelligence*, pp. 1–27, 2010, doi: 10.1007/978-3-642-14435-6_1.

[20] A. Bandi, B. Kongari, R. Naguru, S. Pasnoor, and S. V. Vilipala, “The rise of agentic AI: A review of definitions, frameworks, architectures, applications, evaluation metrics, and challenges,” *Future Internet*, vol. 17, no. 9, pp. 404, 2025, doi: 10.3390/fi17090404.

[21] D. Wu *et al.*, “Generative multi-agent collaboration in embodied AI: A systematic review,” *arXiv preprint*, 2025, doi: 10.48550/arXiv.2502.11518.

[22] D. Grossi, A. Rosenfeld, and N. Talmon, “Advances in multi-agent systems research: EUMAS 2021 Extended selected papers,” *Sn Computer Science*, vol. 4, no. 5, 2023, doi: 10.1007/s42979-023-02098-0.

[23] R. Calegari, G. Ciatto, V. Mascardi, and A. Omicini, “Logic-based technologies for multi-agent systems: A systematic literature review,” *Autonomous Agents and Multi-Agent Systems*, vol. 35, no. 1, 2020, doi: 10.1007/s10458-020-09478-3.

[24] X. Li, S. Wang, S. Zeng, Y. Wu, and Y. Yang, “A survey on LLM-based multi-agent systems: Workflow, infrastructure, and challenges,” *Vicinagearth*, vol. 1, no. 1, 2024, doi: 10.1007/s44336-024-00009-2.

[25] Ebadi, A. Auger, and Y. Gauthier, “Detecting emerging technologies and their evolution using deep learning and weak signal analysis,” *Journal of Informetrics*, vol. 16, no. 4, pp. 101344, Nov 1. 2022, doi: 10.1016/j.joi.2022.101344.

[26] E. Pretel, A. Moya, E. Navarro, V. Lopez-Jaquero, and P. Gonzalez, “Analysing the synergies between multi-agent systems and digital twins: A systematic literature review," *Information and Software Technology*, vol. 174, pp. 107503, Oct. 2024, doi: 10.1016/j.infsof.2024.107503.

[27] Y. Gao *et al.*, “Vagus nerve stimulation paired with rehabilitation for motor function, mental health and activities of daily living after stroke: A systematic review and meta-analysis,” *Journal of Neurology Neurosurgery & Psychiatry*, vol. 94, no. 4, pp. 257–266, 2022, doi: 10.1136/jnnp-2022-329275.

[28] J. A. Ramos-Castaneda, C. F. Barreto-Cortes, D. Losada-Floriano, S. M. Sanabria-Barrera, F. A. Silva-Sieger, and R. G. Garcia, “Efficacy and safety of vagus nerve stimulation on upper limb motor recovery after stroke. A systematic review and meta-analysis,” *Frontiers in Neurology*, vol. 13, 2022, doi: 10.3389/fneur.2022.889953.

[29] X. Wang et al., “Application of vagus nerve stimulation on the rehabilitation of upper limb dysfunction after stroke: A systematic review and meta-analysis,” *Frontiers in Neurology*, vol. 14, 2023, doi: 10.3389/fneur.2023.1189034.

[30] H. Park, and M. Poo, “Neurotrophin regulation of neural circuit development and function,” *Nature Reviews Neuroscience*, vol. 14, no. 1, pp. 7–23, 2012, doi: 10.1038/nrn3379.

[31] G. Leal, C. R. Bramham, and C. B. Duarte, “BDNF and hippocampal synaptic plasticity,” *Vitamins and Hormones*, pp. 153–195, 2017, doi: 10.1016/bs.vh.2016.10.004.

[32] M. A. Cervera et al., “Brain-computer interfaces for post-stroke motor rehabilitation: a meta-analysis,” *Annals of Clinical and Translational Neurology*, vol. 5, no. 5, pp. 651–663, 2018, doi: 10.1002/acn3.544.

[33] L. Zhang, M. Zhang, Y. Zhang, N. Li, J. Hu, and X. Peng, “Efficacy of brain–computer interface with functional electrical stimulation, transcranial direct current stimulation, and conventional therapy on upper limb recovery after stroke: A systematic review and network meta-analysis,” *Frontiers in Neurology*, vol. 16, Art. no. 1643536, 2025, doi: 10.3389/fneur.2025.1643536.

[34] S. Palaniappan, R. Logeswaran, K. Subaramaniam, O. Baker, and B. Ngoc Dung, “Training the brain: A machine learning approach to predicting wellbeing through intentional thought pattern modification”, *Journal of Informatics and Web Engineering*, vol. 4, no. 3, pp. 64–89, Oct. 2025, doi: 10.33093/jiwe.2025.4.3.4.

[35] S. Palaniappan, R. Logeswaran, and Y. L. Yong, “Machine learning model for assessing human well-being using brain wave activities”, *Journal of Informatics and Web Engineering*, vol. 4, no. 2, pp. 93–113, Jun. 2025, doi: 10.33093/jiwe.2025.4.2.7.

BIOGRAPHIES OF AUTHORS

	**Raza Hasan** is a Senior Lecturer in Computing at Solent University, UK. His research focuses on artificial intelligence, machine learning, data mining, data science, and educational data analytics. He also has strong interests in health informatics, learning analytics, and the practical application of data-driven solutions in education and healthcare. He has over 15 years of experience in higher education and has published over 100 research articles in reputable journals and conferences. He can be contacted at email: raza.hasan@solent.ac.uk.
	**Salman Mahmood** is an Associate Professor in the Department of Computer Science at Nazeer Hussain University (NHU), Pakistan. He specializes in artificial intelligence, machine learning, cybersecurity, and data science. He holds academic credentials in computer science and has been actively involved in both teaching and research. Over his career, he has published in areas of AI and related disciplines, contributing to the growth of the research profile at NHU. He can be contacted at email: salman.mahmood@nhu.edu.pk.