Journal of Computer and Knowledge Engineering, Vol. , No.. 2023.
DOI:

A Knowledge Graph and Large Language Model Approach for Mapping Associations Between Cognitive Neuroscience and Marketing*

First A. Sarabadani [1], Second K. Rahsepar Fard [2], and Third M. Tajvidiyan[3]

***Abstract*-- A semi-automated framework to create a structured association map between cognitive neuroscience and marketing, using expert-validated large language models (LLMs), supports a knowledge graph-based approach. The starting point for this framework is an existing cognitive neuroscience knowledge graph (CNKG) of 312 concepts as the source domain and a curated marketing concept list consisting of 100 concepts as the target domain.**

The mapping process consists of two distinct phases: the first phase involves automated association generation via four different LLMs (GPT-4, Gemini, Claude, & Grok) producing candidate associations between cognitive neuroscience and marketing by determining synonymous relationships with their word choices; and the second phase includes human-based expert validation in which marketing and economics experts evaluate, revise, and filter the candidate associations produced by the automated associations.

In total, the automated process generated 1,284 unique candidate relationships. After validation by human experts, there were 742 cross-domain relationships remaining and placed in a bipartite format with a total of 412 nodes and an average degree of connectivity 3.57.

Quantitative analyses indicate that 53.3% of candidate relations were generated from at least 2 LLM's, suggesting moderate inter-model agreement between the four LLM's used to generate candidate relationships. In the next step where each candidate relationship was reviewed by an expert, the expert evaluation showed 39.9% approved the candidate association, 17.9% modified the candidate association, and 42.2% rejected the candidate association. This further emphasizes that human supervision is essential in maintaining the scientific integrity of the mapping process.

In reviewing and analyzing the graph that was produced through this process, many core concepts within neuroscience were identified as being hub nodes (reward processing, emotion, attention, and decision-making). Additionally, there were strong relationships between the hub neuroscience concepts and many marketing constructs (including the brand).

***Keywords:* Cross-domain Knowledge Mapping, Cognitive Neuroscience, Knowledge Graph, Large Language Models, Neuromarketing**

INTRODUCTION

Neuromarketing is a newly emerging field that combines marketing and Cognitive Neuroscience. Neuromarketing focuses on consumer behavior and how our brains create and process decision-making. Despite the rapid growth of this field, we still lack knowledge of the relationship between marketing (brand loyalty, purchase intention, etc.) and the ways in which our brains process the world around us (Attention, Memory, Emotions, etc.) [1]. Researchers and practitioners have a difficult time finding consistent ways to link marketing and the way our brains function in their research.Knowledge Graphs (KGs) can provide accurate and comprehensive representations of nodes, as well as nodes' relationships, in a way that machines can read and understand [2]. They can therefore be used to integrate theoretically sound knowledge into LLMs [2]. Therefore, KGs provide a way to enhance the recommendation capabilities of LLMs by integrating accurate knowledge into their systems [2]. Flexible indexing methods for KGs provide an all-encompassing view of a specific node's neighborhood, therefore enhancing the context in which an LLM stores and processes knowledge, as well as providing for efficient retrieval of relevant knowledge via selective retrieval methods, based on popularity [2].

LLMs have shown that they can extract the meaning of relationships found in unstructured text and make potential connections through known concepts, i.e. building a semantic understanding. They are known to introduce variability (e.g., noise, speculative connections) to the result when they are used alone and therefore need oversight from an expert [2]. Expressive graph encoder approaches have been developed that allow LLMs to utilise structured data when performing work (e.g., overcoming limitations associated with longer context models); this has improved LLM performance for more complex tasks [2]. However, LLMs can still be expensive to build and run and may also become outdated or incomplete over time [2], [3].

Conversely, knowledge graph (KG) technology provides a way to explicitly store information about principles of knowledge using entities and relationships, enabling symbolic

* Manuscript received Revised, , accepted .
[1]Department of Computer Engineering and Information Technology, University of Qom, Qom, Iran, Email:
[2] Corresponding author. Department of Computer Engineering and Information Technology, University of Qom, Qom, Iran, Email:
[3]Department of Computer Engineering and Information Technology, University of Qom, Qom, Iran, Email: www.tajvidiyaan@gmail.com

reasoning along with better recall/access to comparatively rare knowledge. KGs typically have limitations as well; they are often domain-specific, expensive to create, and do not generalise well [2], [3].

To overcome the limitations of LLMs and KGs, Retrieval-Augmented Generation (RAG) systems have been developed, which combine the general design of an LLM with retrieval of data from external knowledge sources. This provides the RAG system with chunks of text to work from when it generates a response ([4]). By using KGs to improve the RAGs with respect to both retrieval precision and reasoning strength, the benefits of using structured knowledge to support understanding and answering of questions involving a higher amount of knowledge can be realised for RAG knowledge-intensive question-answering systems by providing modules that contain reasoning capability, KG integration/validation capability, and validation capability [4], [5]. Also, VKG mapping approaches (e.g., LLM4VKG) provide automation to ontology creation and matching concepts; these approaches have been shown to have advantages over traditional ontology programming/creation methods, while also providing greater robustness in the face of incomplete or very complex ontologies [6]. Finally, the use of Industrial Knowledge Graphs (IKGs) has been shown to improve the semantic organisation of data stored in them, as well as aiding in decision and/or process support within the area of smart manufacturing.

To handle hallucinated output produced by LLM, measures such as BERTScore and BARTScore assess semantic similarity but do not tend to find small differences between the answer spans [8]. Initiatives such as MuShroom-2025 provide more precise identification of hallucinations by identifying exact portions of the hallucinated output thereby enabling finer detection of hallucinations. Integrating KGs with LLM to improve the detection and mitigation of hallucinated outputs will produce more reliable outputs [8]. More recent research shows how Knowledge Graphs are effective at reducing the frequency of hallucinated outputs produced by Large Language Models (LLMs). They include a neurosymbolic framework developed [9], which compares two knowledge graphs: one derived from LLM- generated outputs and another derived from existing reference knowledge bases (or domain knowledge) in order to identify and explain hallucinatory content. Haskins and Adams' results indicate that using graph-based validation methods improves both the detection of factual inconsistencies and the transparency of generated responses. Another study [10] proposed a hybrid Graph Retrieval Augmented Generation (GraphRAG) framework that combines semantic retrieval methods with knowledge graph (KG) reasoning to reduce hallucination rates while improving the reliability of answers and their factual correctness. Both of these studies provide further evidence that hybrid KG-LLM architecture. Security is also an important issue to consider when deploying LLM in an enterprise AI setting. To ensure security, Algorithmic Red Teaming has been used to assess LLM behaviour, with key emphasis on requiring human oversight as well as using hybrid approaches to assessing LLM [11].

In addition, traditional machine learning models have shown that combining predictive algorithms with optimization strategies can improve system robustness and reduce error propagation in complex analytical tasks [12].

Neuromarketing research uses methods taken from neuroscience to provide insights into consumer behaviour and the effectiveness of advertising [13]. For example, EEG recording methods combined with CNN-based channel reconstruction have improved the quality of recording brain activity and enable brain-computer interfaces to be useful in analysing marketing research [13]. The evaluation of large numbers of LLM entails the use of statistical techniques (ANOVA, regression, cluster) to analyse population-level behaviours that are comparable to how cognitive assessments done to humans, thereby providing information about human-like development, capabilities, and training efficiency [14]. Recent developments show how large language models (LLMs) create a powerful platform for connecting marketing and cognitive neuroscience [1]. It's possible for LLMs to reason among many different types of unstructured data, extract objectives at a high level, and create a structured plan from decomposed complex tasks. Therefore, LLMs are helpful for multi-level mapping of cognitive processes and marketing constructs. Multimodal integration contributes to the creation of a richer dataset by combining text-based, numeric, and experimental datasets, thereby increasing the ability to find relationships between neural mechanisms and consumer behaviour. Strategies using structured prompting and hybrid LLM frameworks contribute to reducing variability, increasing consistency, and allowing reproducible, explainable knowledge graph creation [1]. Together, these capabilities make LLMs an important tool for connecting practical marketing applications with theoretical insights in cognitive neuroscience, especially when used with knowledge graphs (KGs) to retrieve and reason.

This research presents a semi-automated methodology which integrates a cognitive neuroscience knowledge graph along with the inferential capabilities of LLMs and expert validation to build a structured association map that connects cognitive neuroscience and marketing. There are four primary contributions of this study:

1. Introduction of a hybrid methodology of integrating knowledge graphs, LLM-based relation extraction, and expert validation.
2. Creation of a bipartite association graph that connects neural concepts to marketing concepts.
3. Providing a quantitative assessment of LLM performance with respect to the influence of expert refinement.
4. Demonstration of the usefulness of the resulting association map for neuromarketing research and practice [1].

TABLE I
Summary of Prior Research on LLM–KG Integration and Neuromarketing Applications

Theme / Concept	Key Points	Applications / Implications	Reference	Relevance to Current Study
Large Language	Extract semantic relations;	Mapping cognitive and behavioral	[1]	Used for inferring associations

Models (LLMs)	generate conceptual associations; multimodal and high-level reasoning	patterns; neuromarketing insights; complex decision-making		between neuroscience and marketing concepts
LLM Challenges	High training cost; knowledge cutoff; hallucinations; output variability	Necessity of hybrid frameworks; expert supervision; structured prompting	[2]	Justifies expert validation and hybrid KG–LLM approach
KG-Augmented LLMs	Multi-hop reasoning; structured plans; multimodal data integration	Mapping neural mechanisms to marketing constructs; explainable, reproducible associations; knowledge-augmented retrieval	[4]	Central methodology of semi-automated association mapping framework
Knowledge Graphs (KGs) & Industrial Knowledge Graphs (IKGs)	Structured, machine-readable entities and relations; addresses semantic fragmentation	Smart manufacturing; IIoT data management; decision-making; cognitive intelligence enhancement	[7]	Foundation for structured representation of neural and marketing concepts
Hallucination Detection in LLMs	BERTScore, BARTScore; span-level detection (MuShroom-2025)	Fine-grained hallucination detection; mitigation strategies; enhanced output trustworthiness	[8]	Ensures reliability of LLM-generated associations
Security in AI & LLMs	Safety, fairness, ethicality; reproducibility challenges	Human-in-the-loop evaluation; Algorithmic Red Teaming; hybrid assessment	[11]	Highlights need for oversight in LLM-assisted frameworks
Neuromarketing	Integration of neuroscience and marketing; EEG and CNN-based EEG signal reconstruction	Understanding consumer behavior; advertising effectiveness; BCIs in marketing research	[13]	Core domain for mapping cognitive processes to marketing constructs
LLM Evaluation			[14]	Supports evaluation
	Population-level analysis; statistical techniques (ANOVA, regression, clustering)	Assessing intrinsic capabilities, developmental patterns, robustness, training efficiency		of LLM outputs and reasoning quality

I. Related work

A. Cognitive Neuroscience Knowledge Graphs

The use of knowledge graphs (KGs) has grown rapidly in the past few years as a means of representing structured relationships among brain regions, cognitive functions, and experimental results in biomedical and neuroscience research [12]. The knowledge represented by these graphs helps to generate hypotheses, integrate various disciplines, and navigate the complexity of knowledge semantically. For example, KGs specific to neuroscience represent the interactions between cognitive functions, brain regions and experimental results, allowing researchers to investigate hidden connections among them.

B. Neuromarketing and Consumer Neuroscience

Marketers have turned to studying the human brain in an effort to understand how viewers react to advertising and products. One such research area is neuromarketing, which utilizes fMRI and psychophysiological measures (e.g., galvanometer) to assess consumer response to various types of marketing stimuli [12]. There are many published studies about how consumers feel emotionally after viewing ads and how those emotions can be linked to their decision-making processes. But there is no true, unified system or framework for relating these neurological responses to marketing construct(s).

KGs can assist in developing a structured body of interdisciplinary knowledge that will connect neural activity, behaviour and marketing.

A thorough review of the literature and a co-citation analysis suggests that there has been a tremendous growth of research conducted at the intersection of neuroscience and business in recent years; however, this growth has mostly occurred in isolated pockets of consumer marketing, management, and psychology, which have very little theoretical connection to each other [17].

More specifically, neuromarketing research has tended to focus primarily on investigating discrete stimulus/response relationships (how people perceive the advertisement) and has not given enough attention to complex, repeated, and socially embedded decision-making processes (i.e., how the person has made this decision over time) that are associated with the marketing and organisations that exist in the world [17].

Co-citation cluster analysis also indicates that much of the current research is based upon a cognitive-oriented, psychological paradigm of thinking, where the researchers do not provide a systematic integration of the underlying physiological brain processes into the development of marketing theory, rather they consider these processes to be a ‘black box’ [17].

These findings highlight a structural gap in the literature, underscoring the need for integrative frameworks capable of continuously bridging neuroscience knowledge with marketing constructs in a coherent and scalable manner [17].

C. Large Language Models for Knowledge Extraction

LLMs (large language models), like GPT-4, Gemini, Claude and Grok, have been used for relation extraction, ontology creation and scientific text mining. LLMs are capable of producing potential relations from textual sources; however, the resulting candidate relations need to be validated due to possible hallucinations or ambiguity in the source domain. Additionally, LLMs are being used to augment datasets, improve semantic understanding and mitigate some of the imbalances found in domain-specific uses, such as reviewer assignment systems.

Because of the fragmented, and interdisciplinary nature of neuroscience based business research, automated processes supporting ongoing, as opposed to one-time, integration of knowledge provide significant benefits in establishing and preserving the consistency of concepts as the source disciplines of the knowledge continue to evolve.

D. Knowledge Graph Construction

Knowledge Graph Construction (KGC) is a process that extracts information from raw text to build a structured representation of knowledge in the form of RDF triples [16]. Early KGC approaches typically used rule-based systems, linguistic-pattern matching techniques, and classical machine learning techniques such as Hidden Markov Models (HMMs), Conditional Random Fields (CRFs), Support Vector Machines (SVMs), etc., for automatic entity and relation extraction. More recently, KGC researchers have begun developing methods utilizing deep learning, with emphasis on Transformer-based networks such as BERT, sometimes combined with Recurrent Neural Networks (RNNs) or Graph Neural Networks (GNNs), for joint entity recognition and relation extraction [16].

There are a variety of benchmark datasets for KGC that cover many application domains. These datasets include: CoNLL04 (general-purpose dataset); SciERC (domain-specific dataset); TransOMSC (domain-specific dataset); DBpedia, Wikipedia, and Wikidata-derived datasets (e.g., WebNLG, FewREL, DocRED, T-REx, TekGen, KELM) (domain-agnostic datasets) [16]. Existing benchmark datasets for knowledge graph construction (KGC) are widely used in both supervised and few-shot learning settings; however, scaling KGC models to large datasets often incurs significant computational cost and resource requirements.

E. LLMs for Knowledge Graph Construction

LLMs have been utilized with limited or no prior training on KGC through their capabilities of using their internal parameters to create RDF triples from text. In contrast, the use of multi-turn question answering or extraction without a schema as seen in the framework of ChatIE and Extract-Define-Canonicalize, is being utilized to create entity and relationship triples. In addition, the integration of LLMs with KGs has shown promise for enhancing Knowledge Graphs (KGs) where the desired output can enhance the density of sparse graphs and thus enhance the overall quality of the graph [12].

F. Code and Small LLMs for KGC

In a recent study, it was shown that Code LLMs are more effective than using natural language prompts for triple extraction because of how they represent triples as class instances in Python as well as employ reasoning techniques (e.g., chain of thought (CoT)) to solve the problem of extracting the triples [12]. Additionally, researchers have been using parameter-efficient ways to fine-tune smaller LLMs (e.g., GPT-2, Falcon 7B, LLaMA 13B) on KGC augmented datasets using methods like LoRA. Therefore, although little computational resources were used in producing the results presented here, the previous literature used mostly to address the problem of performing zero shot/few shot inference or finding new data to augment their training datasets was done using much larger datasets [16].

G. Reviewer Assignment and KGs/LLMs

The reviewer assignment problem - matching papers with their reviewers and ensuring that no conflicts exist is a well-studied area of research that has recently gained traction from the use of KGs and LLMs [12]. Reviewer assignments were conventionally done using either manual techniques or keyword-based methods; both approaches have limitations based on the amount of semantic understanding available and their scalability. Recent attempts have employed various techniques, including topic modelling techniques such as Latent Dirichlet Allocation (LDA), Latent Semantic Indexing (LSI) and neural models such as BERT, CNN, and biLSTM, to improve the degree of alignment between reviewer expertise and accuracy of recommendations made for reviewer assignments. LLMs have provided a new approach to enhance the methods used for data augmentation, semantic embedding generation and few-shot knowledge extraction in order to improve the matching of reviewers to manuscripts [12].

Graph-based methods can be applied to improve fairness, as well as detect potential conflicts-of-interest by utilising co-authorships, institutional affiliations, and distance of collaboration. For example, PeerReview4All and Reviewer Round Robin algorithms provide methods to optimise both coverage and workload balancing, and when combining LLMs with KGs, they provide a means of enriching knowledge required for assignment tasks and provide a greater degree of semantic understanding than is currently available [12].

The combination of KGs and LLMs provides complementary strengths to the assignment of reviewers to manuscripts – KGs contain structured relational knowledge and LLMs contain unstructured semantic knowledge and enable low resource or zero-shot inference. This complementary nature of both viewpoints has led to work on KG construction, domain adaptation, and reviewer assignments recently [12], [16].

II. METHODOLOGY

A. Framework Overview

This research introduces a semi-automated method for discovering connections between cognitive neuroscience and marketing by combining a current CNKG, an enormous number of LLMs, and evaluation by human experts. The procedure includes two sequential steps:

The use of LLMs to generate automated associations (candidate connections) between neuroscience and marketing ideas.

Expert evaluation and modification of the generated associations for the purpose of scientific validity.Inputs:

CNKG: A curated knowledge graph representing cognitive neuroscience concepts and their interrelations [19], [18].

Marketing concepts: A curated list of 100 key marketing concepts from textbooks, expert consultation, and neuromarketing literature.

Outputs:

A validated bipartite knowledge graph encoding cross domain relationships between neuroscience and marketing concepts, enabling systematic knowledge transfer from a specialized scientific domain to an applied domain.

B. Source and Target Concept Preparation

1) Cognitive Neuroscience Concepts (Source Domain)

The extracted concepts were cognitive processes/functions; neural processes/mechanisms; attention; reward processing; memory; emotion; processing of decisions; executive function/control [17,18] (from the original source) along with semantic consistency between all sources, standardized identifiers for all sources, and less ambiguity.

2) Marketing Concepts (Target Domain)

A set of 100 marketing concepts was compiled based on:

- Standard marketing and consumer behavior textbooks.
- Expert consultation in marketing and economics.
- Frequency and relevance in neuromarketing literature.

Examples include brand perception, purchase intention, customer engagement, advertising effectiveness, price sensitivity, and consumer trust.

```
You are an expert in both cognitive neuroscience and marketing.

Consider the following neuroscience concept:
Reward processing

Consider the following marketing concept:
Purchase intention

Identify whether there is a meaningful relationship between these two concepts.

If yes:
- Briefly describe the relationship.
- Provide a clear scientific justification based on cognitive neuroscience or marketing theory.

If no:
- State exactly: "No relationship."
```

Fig. 1. Prompt template for LLM-based knowledge extraction process.

C. Phase 1: Automated Association Generation Using LLMs

Four state-of-the-art LLMs—GPT-4, Gemini, Claude, and Grok—were used as inference engines [17], [18].

1) Prompt Design

Each neuroscience–marketing concept pair was prompted using a structured template, instructing the model to:

- Identify meaningful relationships.
- Provide a concise scientific justification.
- Structured prompting mitigates hallucinations and encourages concept-level reasoning [16].

2) Association Extraction Process

For each neuroscience–marketing concept pair, the LLMs generated candidate associations in the form:

⟨Neuroscience Concept, Marketing Concept, Relationship Description⟩

The outputs from all four models were collected into a unified candidate relation pool.

3) Normalization and Aggregation

Candidate associations were:

- Normalized: synonymous terms unified.
- Deduplicated: repeated associations removed.
- Filtered: trivial or syntactically invalid outputs removed.

This produced a refined set of candidate associations for expert evaluation.

D. Phase 2: Human Expert Validation and Refinement

During this stage, a total of two validators (experts) validated the associations in terms of cross-domain relationships between cognitive neuroscience and marketing. The validators included marketing scholars and data science practitioners with academic and/or professional experience related to the validation of associations. In addition, guidance and consultation were also obtained from cognitive neuroscience experts to support the evaluation process. The validators assessed each association independently of each other based on predetermined evaluation criteria.

1) Evaluation Criteria

Experts assessed associations based on:

- Scientific plausibility – alignment with neuroscience and marketing theories.
- Conceptual clarity – interpretability and well-defined relations.
- Practical relevance – applicability in marketing contexts [15].

2) Labeling Strategy

- Each association was labeled as:
- Accepted: valid and meaningful.
- Modified: valid but requiring refinement.
- Rejected: unsupported or vague.

E. Knowledge Graph Construction and Analysis

The validated associations were encoded as a bipartite knowledge graph:

- One node set: neuroscience concepts.
- Second node set: marketing concepts.
- Edges: validated cross-domain relationships.

Edge annotations:

- Relationship type
- LLM source(s)
- Expert validation label

Graph-theoretic metrics:

Graph density:

$$\text{Density} = \frac{|E|}{|V| \times (|V|-1)} \text{ [25]}$$

These metrics quantify the connectivity and structural properties of the resulting network.

F. Reliability and Quality Control

Inter-model Agreement Rate (IAR):

$$\text{IAR} = \frac{\text{Number of agreements between models}}{\text{Total number of assignments}} \times 100$$

Indicates the consistency of LLMs in identifying meaningful conceptual associations [15].

Expert modification/rejection rate: quantifies LLM noise and validates the necessity of human-in-the-loop curation.

G. Reproducibility

All prompts, concept lists, and evaluation criteria are fully documented, ensuring reproducibility. The framework is domain-independent and can be adapted for other interdisciplinary knowledge mappings.

Pseudocode Summary

```
# 1. Automated association generation
candidates = []
for nc in CNKG:
    for mc in Marketing:
        for model in LLMs:
            rel =
model.predict_relation(nc, mc)
            candidates.append(rel)
candidates = deduplicate(candidates)

# 2. Expert validation
validated = []
for rel in candidates:
    label = expert_evaluate(rel)   #
Accept, Modify, Reject
    if label != "Reject":
        validated.append(refine(rel,
label))

# 3. Graph construction
KG = create_bipartite_graph(validated)
# 4. Evaluation
metrics = compute_metrics(candidates,
validated, KG)
Output: KG, metrics
```

III. Experimental Setup

A. Data Sources

The two main datasets were chosen to study the ability of longitudinal learning mechanics to create relationships through automated knowledge extraction between knowledge domains that use different ontological structures. Accordingly, the data sources are:

1. The Cognitive Neuroscience Knowledge Graph (CNKG): This is an existing Cognition related knowledge graph which has been developed using the peer reviewed cognitive neuroscience literature. The CNKG consists of 312 neuroscience concepts and related to Cognition, emotion, perception and decision-making. The CNKG includes each of the neuroscience entities (e.g. concepts) linking together within the broad scientific corpus of knowledge and serves as the source domain for cross-domain mapping - [19].

2. Marketing Concept Set: The Marketing Concept Set consists of a collection of 100 available marketing concepts that have been compiled from the body of established marketing and consumer behaviour textbooks and refined by experts. This dataset represents the target domain and consists of fundamental constructs including consumer segmentation, brand equity, purchase intention, and advertising effectiveness [20], [21], [22], [23].

The two datasets above were selected because it was thought that an examination of how LLMs have the ability to derive meaningful analogy and/or relational-based relationships between the two domains (knowledge domains) will identify patterns that will help suggest that cross-domain knowledge association extraction can be achieved.

B. Knowledge Extraction Workflow

The experimental process was completed in the following steps:

1. Candidate Relation Generation : Multiple large language models (LLMs) were used to generate candidate cross-domain relationships between cognitive neuroscience and marketing concepts. This was done using structured prompts constructed from CNKG and the Marketing Concept Set, following recent approaches in LLM-based knowledge graph construction and relation extraction [24], [25].

2. Relation Filtering and Normalization : Candidate relations were consolidated into a standardized set, where duplicate and semantically equivalent relations were removed and unified to ensure consistency and uniqueness. This follows established preprocessing and normalization practices in knowledge graph construction and LLM-based information extraction pipelines [25].

3. Expert Review : Cognitive neuroscience and marketing experts independently evaluated each candidate relation based on semantic plausibility, conceptual clarity, and relevance. Each relation was categorized as Accepted, Modified, or Rejected, consistent with human-in-the-loop validation

frameworks for LLM-generated knowledge structures [25].

4. Graph Integration : The validated relations were integrated into a bipartite knowledge graph for further structural and topological analysis of cross-domain relationships. This aligns with recent methodologies in LLM-augmented knowledge graph construction and graph-based reasoning systems [24], [26].

C. *Evaluation Metrics*

The following performance metrics were determined to quantitatively analyze how well LLMs are performing and the impact of expert validation:

- Inter Model Agreement Rate (IAR): Percentage of relations that at least two different LLMs suggested that showed consensus,
- Expert Acceptance Rate (EAR): Percentage of LLM suggested relations accepted without any modification by experts,
- Modification Rate (MR): Percentage of the relations generated by LLMs that were modified by experts,
- Rejection Rate (RR): Percentage of the relations generated by LLMs that were rejected by experts.

In addition to these human-centric performance metrics, other graph structure metrics were calculated to assess the effect of integration on the expanded knowledge graph which included:

- The number of nodes and edges (which show basic size expansion of the graph)
- The mean number of connections per entity (the average node degree),
- The ratio of actual edges to all possible edges in the graph (the density of the graph).

These structural indicators are well established metrics for determining connectivity and structural robustness within a KG as well as how well new relations integrate into the topology of the graph.

D. *Robust Experimental Design*

The validity and reproducibility of results in the experimental design were addressed through the use of several controls:

- Use of multiple LLM Backbones: The use of a minimum of two different LLM architectures to lessen the chance of bias due to the eccentricities of any one model.
- Cross Validation: A set of proposed relations for a small set of CNKG entities were withheld from the construction of the overall dataset and later evaluated independently in order to measure LLM's generalizability.
- Blind Review Process: Expert reviewers evaluated proposed relations in random order and without knowledge of the specific model used to propose them.

This design permitted systematic measurement of automated relation generation capability, validity of human expert validation, and relationship structure and graphical evolution/changes as a function of LLM incorporation of derived associations across domains.

IV. Results

A. *Overall Knowledge Graph Statistics*

The automated LLM phase generated a total of 1,284 candidate associations between neuroscience and marketing concepts.

After expert validation, the final bipartite knowledge graph contained:

- 312 neuroscience nodes (from CNKG)
- 100 marketing nodes

742 validated cross-domain relationships (edges)

The detailed statistics of the final knowledge graph are summarized in the table below**:**

TABLE II
Presents The Counts Of Nodes And Validated Relationships In The Knowledge Graph

Feature	Count
Candidate associations generated by LLMs	1,284
Neuroscience nodes (from CNKG)	312
Marketing nodes	100
Validated cross-domain relationships (edges)	742

The corresponding chart provides a visual summary of the knowledge graph statistics:

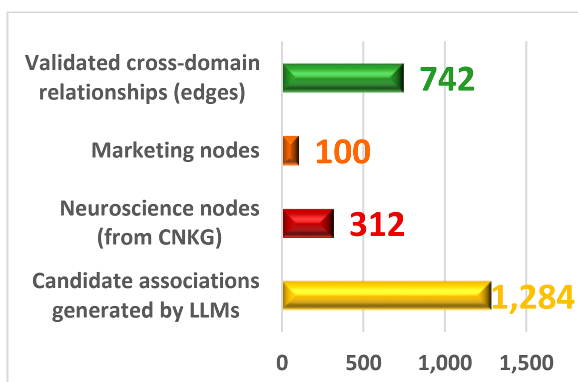

Fig. 2. illustrates the data presented in Table 1 in graphical form.

1) *Graph Metrics:*

To quantify the structural properties of the constructed knowledge graph, two standard graph metrics were computed: graph density and average node degree.

2) Graph Density:

Graph density measures the proportion of actual edges present relative to the maximum possible number of edges in the graph. For a directed graph, it is defined as:

$$\text{Density} = \frac{|E|}{|V|\times(|V|-1)} \text{ [24]}$$

where $|E|$ is the number of edges and $|V|$ is the number of nodes.
In our case:

$$\text{Density} = \frac{742}{100 \times 312} = 0.0238$$

This relatively low density indicates that the graph is sparsely connected, which is typical for knowledge graphs spanning multiple domains.

The average node degree was **3.57**, indicating moderate connectivity between the two domains.

3) Average Node Degree:

The average node degree provides the mean number of edges connected to each node:

$$\text{Average Degree} = \frac{2 \times |E|}{|V|}$$

For our graph, this yields:

$$\text{Average Degree} = 3.57$$

indicating moderate connectivity between the two domains.

According to these findings, the structural connections between the various nodes indicate moderate levels of connectivity and characterize the features of the built knowledge graph, and while there are primarily sparse connections for most of the nodes, there are enough connections to support typical tasks at a later stage.

B. LLM Performance Analysis

Among the 1,284 candidate relations:

- 684 relations (53.3%) were proposed independently by at least two LLMs.
- 241 relations (18.8%) were suggested by all four models.
- 197 relations (15.3%) were suggested by three models.
- 246 relations (19.2%) were suggested by two models.
- 600 relations (46.7%) were unique to a single model.

The detailed distribution of LLM-suggested relations is summarized in the table below:

TABLE III
Presents The Breakdown Of Relations Proposed By Different Numbers Of LLMs.

Number of LLMs Suggesting Relation	Number of Relations	Percentage
All four models	241	18.8%
Three models	197	15.3%
Two models	246	19.2%
At least two models (total)	684	53.3%
Unique to a single model	600	46.7%
Total	1.284	**100%**

The corresponding chart visualizes the distribution of relations proposed by the LLMs:

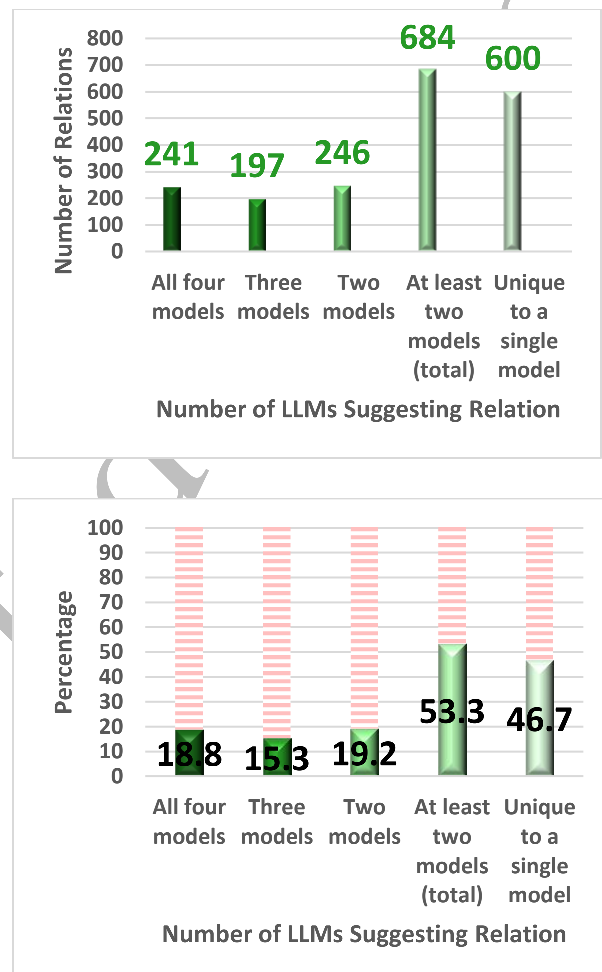

Fig. 3. provides a graphical representation of the data summarized in TABLE III

Inter-Model Agreement Rate (IAR): To quantify the agreement between multiple models in assigning reviewers, we define the Inter-Model Agreement Rate (IAR) as the percentage of reviewer assignments for which all considered models agree.
Formally:

$$\text{IAR} = \frac{\text{Number of agreements between models}}{\text{Total number of assignments}} \times 100$$

Where:

- *Number of agreements between models* is the count of assignments for which all models recommend the same reviewer for a manuscript.

- *Total number of assignments* is the total number of manuscripts assigned across all models.

For our experimental setup, the IAR was calculated as:

IAR = 53.3%

This suggests that 53.3% of all manuscript project submissions were consistent across all evaluated models.

This indicates areas of both agreement and disagreement among these evaluated models.

The results show a moderate degree of consistency when evaluating LLMs in making conceptual associations.

C. *Impact of Expert Validation*

Expert evaluation of the 1,284 candidate relations yielded the following outcomes:

- 512 relations (39.9%) were accepted without modification.
- 230 relations (17.9%) were modified to improve scientific clarity and precision.
- 542 relations (42.2%) were rejected.

The detailed outcomes of expert validation are summarized in the table below:

TABLE IV

Presents A Breakdown Of Accepted, Modified, And Rejected Relations.

Expert Evaluation Outcome	Number of Relations	Percentage
Accepted without modification	512	39.9%
Modified for clarity/precision	230	17.9%
Rejected	542	42.2%
Total	1.284	**100%**

The corresponding chart visualizes the expert validation outcomes:

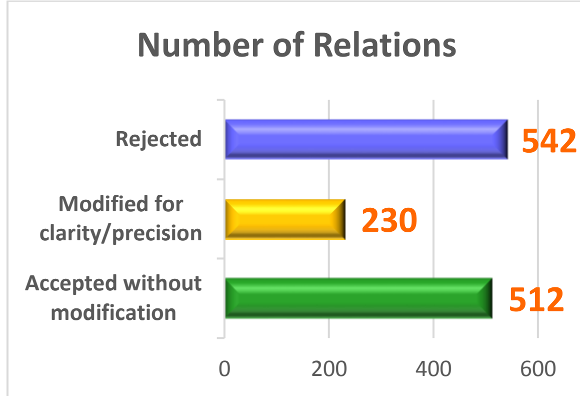

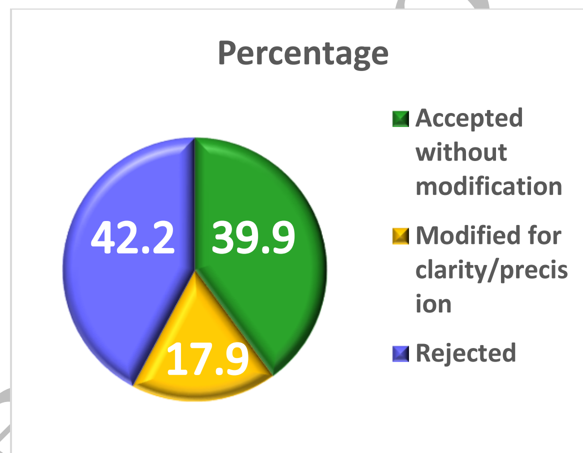

Fig.4. provides a graphical representation of the data summarized in TABLE IV.

The main causes of rejection were:

1. Conceptual vagueness (38% of rejected relations),
2. Lack of neuroscientific evidence (34%),
3. Overly generic associations (28%).

The detailed breakdown of these rejection causes is presented in the table below:

TABLE V

Summarizes the Main Reasons for Relation Rejection

Cause of Rejection	Percentage of Rejected Relations
Conceptual vagueness	38%
Lack of neuroscientific evidence	34%
Overly generic associations	28%

The corresponding chart visualizes the distribution of rejection causes:

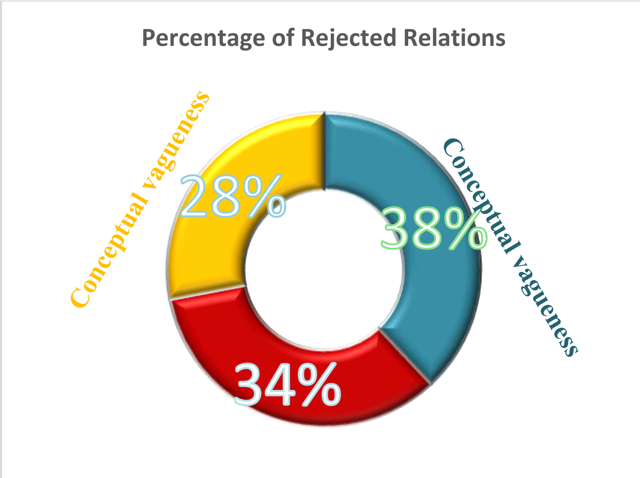

Fig.5. provides a graphical representation of the data summarized in TABLE V.

This confirms the critical role of human supervision in refining LLM-generated knowledge.

D. Distribution of Relationship Categories

The thematic categorization of the validated relations (742 edges) resulted in the following grouped relations:

- 214 emotion-based relations, 28.8% were related to emotion and loyalty of the brand and to the recall of advertisement materials.
- The theme of reward and motivation consisted of 183 relations, 24.7% percentage and examples of reward processes related to purchase intentions.
- The theme of attention and perception contained 162 relations and comprised 21.8%. Examples include the relation of attention to advertisement effectiveness.
- The working memory and brand recognition were included in the theme of memory and comprised 121 relations, 16.3%.
- The theme of decision and executive control included 62 relations, constituting 8.4%.

The above results show that emotional and reward-based neural processes dominate marketing-related phenomena.

The detailed counts for each relationship category are summarized in the table below:

Table VI
Presents the Full Distribution of Validated Relations Across These Thematic Categories

Relation Category	Number of Relations	Percentage of Total	Examples
Emotion-related associations	214	28.8%	Emotion → brand loyalty, emotion → advertising recall
Reward & motivation mechanisms	183	24.7%	reward processing → purchase intention
Attention & perception mechanisms	162	21.8%	Attention → advertising effectiveness
Memory-related mechanisms	121	16.3%	working memory → brand recognition
Decision-making & executive control	62	8.4%	–

The corresponding chart visualizes the distribution of these relationship categories:

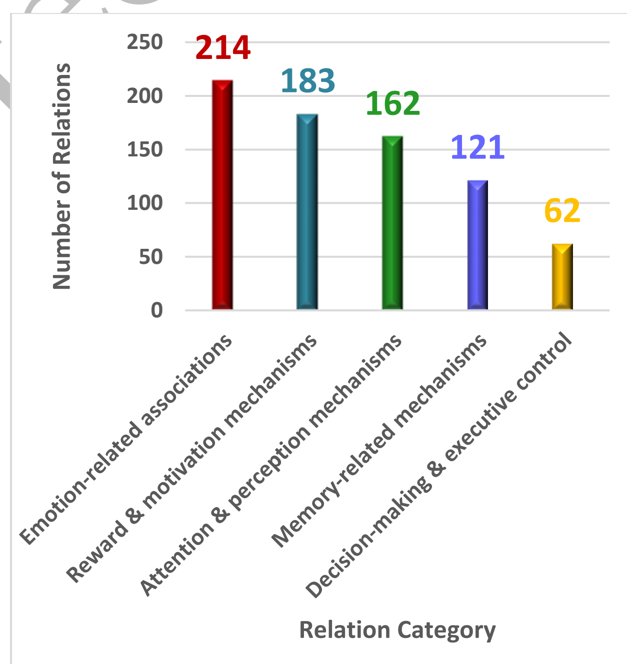

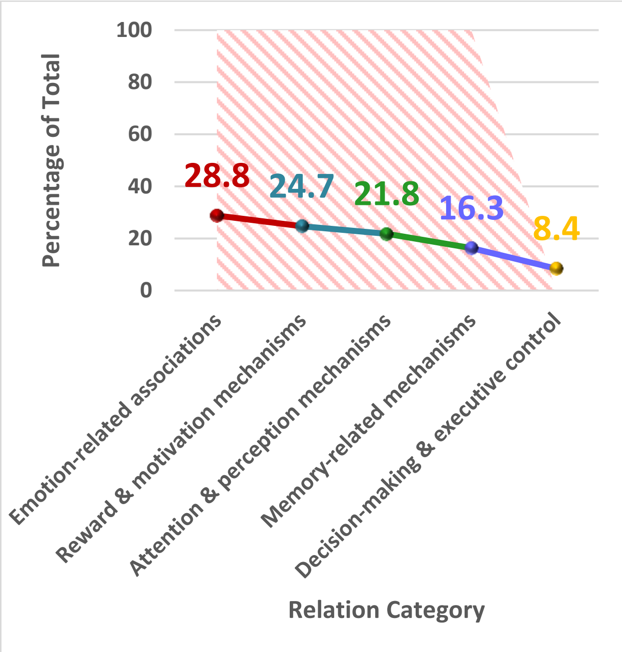

Fig. 6. provides a graphical representation of the data summarized in Table VI.

E. *Graph Structural Analysis*

Centrality analysis revealed that a small number of neuroscience nodes acted as hubs in the bipartite graph:

Table VII
Lists the Centrality Values for Each Node

Neuroscience Concept	Degree
Reward Processing	41
Emotion	38
Attention	35
Decision-Making	29
Memory	26

Figure 6 provides a graphical representation of the centrality scores listed in the Table VII:

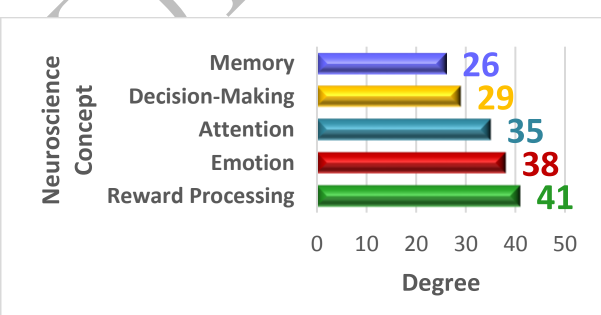

Fig.7. Graphical representation of the centrality scores for neuroscience nodes in the bipartite graph, highlighting hub nodes identified in the analysis.

These hub nodes were strongly connected to marketing concepts such as:

- brand loyalty
- advertising effectiveness
- consumer trust
- purchase intention
- customer engagement

This hub structure suggests that a limited set of neural mechanisms accounts for a large portion of marketing behavior explanations.

F. *Reliability Indicators*

There are two quantitative indicators that indicate that the proposed framework is valid:

1. The inter-model agreement rate of 53.3% implies that there is agreement among the various LLMs regarding a consistent subset of conceptual relationships.
2. The expert rejection and modification rate of 60.1% emphasizes the value of using experts in the analysis of data to ensure that the scientific quality of any conclusions drawn from the data is of high standard.

The detailed values for these indicators are summarized in the table below:

Table VIII
Presents a Summary of These Quantitative Reliability Indicators

Indicator	Value	Interpretation
Inter-model agreement rate	53.3%	Shows convergence among different LLMs on a stable subset of relations
Expert rejection/modification rate	60.1%	Highlights the importance of human-in-the-loop validation to ensure scientific rigor

Together, these results confirm that combining LLM-based inference with expert validation provides a robust methodology for constructing cross-domain knowledge graphs.

V. Discussion

According to the literature reviews and co-citation analysis, business and management have expressed interest in using neuroscience for future research; whereas less is known about its use in B2B marketing. The exploration of this area through the combined lens of neuroscience and B2B marketing yielded only a few examples: nine out of nine existing studies focused on a one time exchange (i.e., Type II exchanges) as opposed to either decision-making processes or frequent complex exchanges with substantial value offerings within B2B. As such, the opportunity to increase knowledge through incorporating findings from neuroscience is quite large. While there is somewhat of a lack of research linking neuroscience to B2B marketing, this is not the primary reason for providing additional opportunities for the two fields; instead, the decision to incorporate findings from neuroscience because there is clearly a need and an opportunity to do so.

The findings from the literature review and co-citation analysis in business and management with regard to the use of neuroscience indicate neuroscience could act as an effective methodology both as a second method of research for social science and natural science as well as a foundation for combining social science and natural science [17].

B2B marketing research, using neuroscience, can help to answer both 'how' and 'why' questions, such as those related to the behaviours of social beings, human-machine behaviour, decision making, and the interaction between humans and machines. Some of the areas in which neuroscience may help B2B marketers includes the digitalisation of businesses, the implementation of artificial intelligence, the growing importance of multi-agent systems (that is, systems which rely on the coordinated activities of multiple agents), and innovative ways of thinking about the future. The goal of B2B marketers using neuroscience in research is to link the cognitive or neural aspects of decision behaviours to the knowledge B2B marketers already have about how firms create value - and about their repeated interactions or exchanges with one another.

Social or organisational neuroscience will create new and improved ways for B2B marketers to think about multi-agent systems, as B2B marketing offers the inter-organisational aspect of multi-agent systems. The study of physiological brain activity offers other viable research opportunities for B2B marketers. For example, the use of physiological brain activity may be useful in explaining how managers use forecasting and prediction (the psychological process involved in dealing with future events - forecasting) differently from managers dealing with past events (the psychological process involved in dealing with past events - forecasting).

Given the growing industry awareness that non-marketplace logics can supplement the traditional logic of profit maximisation in the marketplace, the application of neuroscience to B2B marketing will extend the area of cognition to themes embedded in their recent research on values and repeat exchanges.

Multiple opportunities to employ neuroscience insights to expand B2B marketing bridges with neuroscience exist due to the vast subject matter associated with research derived from physiological brain activity to date, which go far beyond B2B marketing's current limited linking of neuroscience with discrete exchanges.

What is the appropriate method for integrating both B2B Marketing and Neuroscience based research? Cluster analysis suggests that business and management research does not have a defined level of coherence between itself and Neuroscience. This provides two options for researching the two disciplines - developing future research based on within discipline advancements (an internalized approach), or continuously integrating both disciplines (a continuous integration approach).

An internalized research approach allows researchers to accumulate knowledge relevant to their field, however it runs the risk of being obsolete or irrelevant due to the rapid progression of scientific and neurobiological advancements.

A continuous integration research approach keeps B2B Marketing researchers engaged with Neuroscience developments while providing a level of divergence between B2B Marketing and Neuroscience based research. Therefore, an initial continuous interdisciplinary integration between B2B Marketing and Neuroscience research will be an optimal model and will potentially lead to a dual research approach in B2B Marketing that integrates both within discipline advancements along with an ongoing analysis of Neuroscience based research advancements.

VI. Applications

There are many possible practical and scientific applications for the proposed Association Map, including:

- Hypothesis Generation in Neuromarketing Research The Knowledge Graph can be used to systematically derive new research questions regarding the neural basis of consumer behaviour and marketing effectiveness.
- Education and Training Tools for Interdisciplinary Learning The Association Map can serve as a structured resource for teaching and learning at the intersections of Cognitive Neuroscience, Marketing and AI.
- Decision Support Systems for Marketers The framework can connect a marketer's marketing strategies to the underlying cognitive/emotional mechanisms used by consumers; thus, supporting practitioners in creating more effective psychologically grounded campaigns.
- AI-Driven Recommendation and Analytics Systems The Knowledge Graph can be used as the Knowledge Architecture for intelligent Systems to integrate insights from Consumer Neuroscience into Personalized Recommendation Engines and Marketing Analytics Platforms.

VII. Limitations and Future Work

Although this study's findings appear encouraging, it has many limitations. The primary limitation is that the quality of association maps and their relationship to Cognitive Neuroscience Knowledge Graph (CNKG) data are dependent on the original data available. Therefore, if the original data contains too little or has biases present, it could result in the inaccuracy of the entire cross-domain association mapping process.

Although LLMs provide strong inferences based on existing data, the associations generated by LLMs contain latent biases or oversimplifications in their training dataset. This suggests a need for validation processes by human experts to minimize errors resulting from LLM processing.

Lastly, an additional limitation of our study is that there was a limited number of expert validators; therefore, this limitation could have an effect on generalizability of these validation findings across multiple disciplines.

In future studies, we plan to explore the following:

- Expanding our framework to include other areas such as psychology, behavioural economics and social neuroscience.
- To create automated systems to score the degree of confidence associated with created associations.
- To match the empirical data from neuroimaging studies and experiments with the knowledge graph in order to ground the graph to observed neural data.
- To create dynamic and temporal representations of the graph to model changes in consumer cognition over time.

VIII. Conclusion

We have developed a semi-automated framework that integrates knowledge graphs, LLMs, and human experts to create an organized association map between cognitive neuroscience and marketing concepts.

By combining a curated cognitive neuroscience knowledge graph with LLMs for inferences, along with the expertise of a human expert, we can systematically discover and confirm cross-domain associations that would be otherwise impossible to glean from the fragmented literature on these topics.

While our tests demonstrated that LLMs can generate many valuable candidate associations (53.3% inter-model agreement), the results suggest that over 60% of those associations needed to be changed or discarded by experts; this reinforces the need for human intervention throughout the process to ensure the validity and clarity of scientific concepts.

By combining a curated cognitive neuroscience knowledge graph with LLMs for inferences, along with the expertise of a human expert, we can systematically discover and confirm cross-domain associations that would be otherwise impossible to glean from the fragmented literature on these topics.

While our tests demonstrated that LLMs can generate many valuable candidate associations (53.3% inter-model agreement), the results suggest that over 60% of those associations needed to be changed or discarded by experts; this reinforces the need for human intervention throughout the process to ensure the validity and clarity of scientific concepts.

The final result was a bipartite knowledge graph indicating how various neural mechanisms (e.g., reward-processing, emotion, attention and decision-making) can be used to explain different marketing concepts (i.e., brand loyalty, purchase intention, and advertising effectiveness). The scope of the bipartite knowledge graph offers both theoretical insight into the neural basis for consumer decision-making, as well as a practical tool for advancing neuromarketing research.

This project is exploratory in that it has produced insights into how hybrid AI/human methods may provide an effective and efficient way to create interdisciplinary knowledge bridges across scientific domains (e.g., cognitive neuroscience and marketing) and within applied fields (e.g., neuromarketing).

The hybrid AI/human approach also contributes to the theoretical development of neuromarketing while providing a basis for the future creation of intelligent decision-support systems and data-supported marketing strategies based on the science of cognitive neuroscience.

IX. Limitations and Future Work

The study has several limitations. First, the quality of the association map heavily depends on the original Cognitive Neuroscience Knowledge Graph (CNKG) data; any biases or insufficient data in the source domain may negatively affect the entire cross-domain mapping process. Second, although large language models provide powerful inferences, their generated associations may contain latent biases or oversimplifications inherited from their training datasets, which necessitates expert validation to minimize errors. Third, the limited number of human expert validators in this study may affect the generalizability of the validation findings across multiple disciplines.

For future work, the authors plan to expand the framework to include additional domains such as psychology, behavioral economics, and social neuroscience. They also aim to develop automated systems to assign confidence scores to each generated association, integrate empirical data from neuroimaging studies and experiments to ground the knowledge graph in observed neural data, and create dynamic and temporal representations of the graph to model changes in consumer cognition over time.

References

[1] F. Aghaee and H. R. Shaker, "RB-LLM control: An intelligent control framework with rule-based large language model decision-making," *Aerospace Science and Technology*, vol. 168, p. 111259, 2025.

[2] S. Wang, W. Fan, Y. Feng, S. Lin, X. Ma, S. Wang, and D. Yin, "Knowledge Graph Retrieval-Augmented Generation for LLM-based Recommendation," in *Proc. 63rd Annu. Meeting Assoc. Comput. Linguistics (ACL)*, 2025, pp. 27152–27168.

[3] C. Voinea, S. Porsdam Mann, J. Savulescu, and B. D. Earp, "The

calculator analogy: Epistemic virtues for using LLMs," *Technology in Society*, vol. 85, p. 103198, 2025.

[4] J. Linders and J. M. Tomczak, "Knowledge graph-extended retrieval augmented generation for question answering," *Applied Intelligence*, vol. 55, no. 17, Art. no. 1102, 2025.

[5] C. Ma, Y. Chen, T. Wu, A. Khan, and H. Wang, "Large language models meet knowledge graphs for question answering: Synthesis and opportunities," in *Proc. Conf. Empirical Methods Natural Language Processing (EMNLP)*, 2025, pp. 24589–24608.

[6] G. Xiao, L. Ren, G. Qi, H. Xue, M. Di Panfilo, and D. Lanti, "LLM4VKG: Leveraging large language models for virtual knowledge graph construction," in *Proc. 34th Int. Joint Conf. Artificial Intelligence (IJCAI)*, 2025, pp. 4715–4723.

[7] X. Li, P. Zheng, J. Bao, L. Gao, and X. Xu, "Achieving cognitive mass personalization via the Self-X cognitive manufacturing network: An industrial knowledge graph- and graph embedding-enabled pathway," *Computers & Industrial Engineering*, vol. 180, p. 109709, 2025.

[8] E. Lavrinovics, R. Biswas, J. Bjerva, and K. Hose, "Knowledge graphs, large language models, and hallucinations: An NLP perspective," *Journal of Web Semantics*, vol. 85, Art. no. 100844, 2025.

[9] R. Haskins and B. Adams, "KEA Explain: A neurosymbolic framework for explaining large language model hallucinations through knowledge graph validation," Journal of Intelligent Systems, vol. 5, no. 1, 2026, doi: 10.1177/29498732261443099.

[10] J. Liu and J. Tang, "LeAK-GraphRAG: A legal academic knowledge GraphRAG framework with hybrid retrieval for mitigating hallucinations in large language models," in Frontiers in Artificial Intelligence and Applications, vol. 411, Amsterdam, The Netherlands: IOS Press, 2026, pp. 597–604, doi: 10.3233/FAIA251624.

[11] S. Jauhari, "Algorithmic red teaming approaches to secure LLMs," *Machine Learning with Applications*, vol. 23, p. 100815, 2025.

[12] Chehreh, S., & Sarabadani, A. (2024). A model based on random forest algorithm and Jaya optimization to predict bank customer churn. Journal of Engineering Management and Soft Computing, 9(2), 132-148. doi: 10.22091/jemsc.2024.9541.1174.

[13] M. Q. Pérez, S. López Bernal, E. Horna Prat, L. Montesano Del Campo, L. Fernández Maimó, and A. Huertas Celdrán, "EEG channel reconstruction using convolutional neural networks in limited BCIs: A proposed method for neuromarketing applications," *Applied Soft Computing*, vol. 181, p. 113455, 2025.

[14] K. Sun and R. Wang, "Breaking myths in LLM scaling and emergent abilities with a comprehensive statistical analysis," *Neural Networks*, vol. 183, p. 108164, 2025.

[15] F. Bagheri, D. Buscaldi, and D. Reforgiato Recupero, "Leveraging knowledge graphs and LLMs for content based reviewer assignment," *Journal of Intelligent Information Systems*, 2025.

[16] P. Gajo and A. Barrón-Cedeño, "Natural vs programming language in LLM knowledge graph construction," *Information Processing & Management*, vol. 62, p. 104195, 2025.

[17] C. Öberg, "Neuroscience in business-to-business marketing research: A literature review, co-citation analysis and research agenda," *Industrial Marketing Management*, vol. 96, pp. 204–217, 2021.

[18] J. Wei, Y. Tay, R. Bommasani, X. Wang, C. Raffel, B. Zoph, S. Borgeaud, D. Hendrycks, and others, "Emergent abilities of large language models," *arXiv preprint*, 2022.

[19] N. Ibrahim, S. Aboulela, and A. Ibrahim, "A survey on augmenting knowledge graphs (KGs) with large language models (LLMs)," *Knowledge and Information Systems*, 2024.

[20] H. Purnomo and Y. Novita, "The impact of branding, advertising, and social media marketing on consumer purchase intention, brand loyalty, and brand equity in the Indonesian market," *Journal of Management and Bussines (JOMB)*, vol. 6, no. 3, Art. no. 9653, 2024.

[21] M. M. Gutiérrez Gutiérrez, J. J. Perona Páez, and F. D. P. Gutiérrez Bonilla, "Models of brand equity: A systematic and critical review," *Cogent Business & Management*, vol. 11, no. 1, Art. no. 2433168, 2024.

[22] L. D. Gonen, T. Tavor, and U. Spiegel, "Unlocking market potential: Strategic consumer segmentation and dynamic pricing for balancing loyalty and deal seeking," *Mathematics*, vol. 12, no. 21, Art. no. 3364, 2024.

[23] E. Tarta and L. Pasaribu, "The impact perceived quality, advertising effectiveness, social media presence and brand image on purchase intention Telkomsel provider," *Enrichment: Journal of Management*, vol. 14, no. 4, pp. 694–704, 2024.

[24] A. Sarabadani and K. R. Fard, "GDPKG-LLM: Integrating Gene, Disease, and Pharmacogenomics Knowledge Graphs for Cognitive Neuroscience Using Large Language Models," *Computational Science and Information*, 2025. https://doi.org/10.7494/csci.2025.26.3.6673

[25] Ibrahim, N., Aboulela, S., Ibrahim, A., & Kashef, R. (2024). "A survey on augmenting knowledge graphs (KGs) with large language models (LLMs): Models, evaluation metrics, benchmarks, and challenges," *Discover Artificial Intelligence*, 4, 76.

[26] S. Seo, H. Cheon, H. Kim, and D. Hyun, "Structural Quality Metrics to Evaluate Knowledge Graph Quality," *Semantic Web Journal*, 2023.