ktu faculty of social sciences, arts and humanities 1922

48/2026
Research Journal
Studies about Languages
pp. 5-19
ISSN 1648-2824 (print)
ISSN 2029-7203 (online)
DOI 10.5755/j01.sal.1.48.43076

LINGUISTICS / KALBOTYRA

Brain metaphors generated by humans and large language models: Evidence from metaphor elicitation

Received 10/2025 | Accepted 05/2026

HOW TO CITE: Urbonaitė, J. & Shcherbino, Y. (2026). Brain metaphors generated by humans and large language models: Evidence from metaphor elicitation. *Studies about Languages / Kalbų studijos*, 48, 5–19. https://doi.org/10.5755/j01.sal.1.48.43076

Brain metaphors generated by humans and large language models: Evidence from metaphor elicitation

Žmonių ir didžiųjų kalbos modelių generuojamos smegenų metaforos: metaforų išgavimo tyrimas

JUSTINA URBONAITĖ, Vilnius University, Lithuania

YULIA SHCHERBINO, Vilnius University, Lithuania

Abstract

Metaphor production constitutes a useful analytical framework for examining how humans and large language models (LLMs) structure meaning beyond literal description. This study compares metaphorical descriptions of the brain elicited from 61 human participants and eight publicly available LLM-based chat systems (40 AI-generated responses). Using matched elicitation prompts (e.g., “The brain is like…”), metaphorical lexis was identified following the MIPVU procedure, and responses were categorised according to their underlying source domains. Both datasets demonstrate a marked tendency toward entrenched technological metaphor patterns, most notably THE BRAIN IS A COMPUTER metaphor. However, human responses draw on a broader range of everyday and embodied source domains (e.g., brain as a muscle, sponge, or knitting yarn), whereas LLM outputs predominantly employ more abstract, system-oriented analogies (e.g., brain as a network, city, control centre) and more frequently render cross-domain mappings explicit. AI-generated responses are also substantially longer, averaging approximately five times the length of human responses. Under these elicitation conditions, LLMs reproduce conventional explanatory metaphors but exhibit reduced diversity in source domain variation and fewer mappings grounded in experiential domains. The results indicate both the capacities and the constraints of LLMs in metaphorical “reasoning,” particularly their reliance on conventionalised patterns and their limited embodied metaphorical grounding.

KEYWORDS: brain metaphors, figurative language, large language models (LLMs), metaphor elicitation, metaphorical competence, MIPVU.

Introduction

The linguistic competence of artificial intelligence (AI), and in particular LLMs, has become a prominent subject of inquiry in linguistics, cognitive science, and AI research. Although LLMs generate fluent language and can produce coherent extended text, the extent to which their outputs simulate human communication beyond surface well-formedness remains an open empirical question (Chang & Bergen, 2024). Recent scholarly work suggests that, while LLMs often match normative patterns in written language, they may diverge from human performance in interactional and affective dimensions, for example, in modelling conversational dynamics learned from dialogue (Umair et al., 2024) or in approximating stylistic and emotional variation typical of human-authored texts (Muñoz-Ortiz et al., 2024). Such findings motivate targeted linguistic tests that go beyond grammaticality and topical coherence and instead examine how meaning is structured and motivated in use.

Figurative language provides a useful basis for investigating this issue. Metaphor, in particular, links linguistic expression to conceptual organisation: it enables speakers to structure abstract conceptual domains in terms of more concrete ones (Lakoff & Johnson, 2003). If metaphors reflect habitual ways of construing complex target domains, then patterns in metaphor production can be informative about what resources a producer draws on when asked to explain or reframe a domain.

Methodologically, metaphor elicitation offers a controlled route into this question. Unlike discourse-based metaphor research, which infers metaphor use from naturally occurring texts, elicitation tasks ask participants to produce metaphors directly (Wan & Low, 2015). This design supports direct comparison across groups because prompts can be matched closely, and responses can be analysed with established procedures such as MIPVU (Steen et al. 2010). For the present study, elicitation was used to compare metaphor production by humans and LLM-based systems under parallel task conditions.

To maintain a clear focus and target domain specificity, the present study focuses on metaphors for the brain. Brain metaphors are historically prominent in scientific and popular discourse and have been shown to shape explanatory habits, not least through persistent computational comparisons (Draaisma, 2000; Dupuy, 2009; Brette, 2019). Brain metaphors are therefore both familiar enough to elicit from non-specialists and theoretically diagnostic for examining whether LLM outputs reflect or compress the variety present in human conceptualisation.

The aim of this study is to provide an empirically grounded comparison of human- versus LLM-generated brain metaphors using a matched elicitation design and MIPVU-guided analysis. The study addresses three research questions:

1. Which source domains are the most frequently evoked by humans and by LLMs when asked to metaphorically describe the brain?
2. How do the datasets compare in terms of their source domain ranges and diversity?
3. How do humans and LLMs differ in the extent to which they explicitly articulate cross-domain mappings?

Theoretical Foundations and Review of Previous Research

Metaphor is commonly defined as "the phenomenon whereby we talk, and possibly think, about one thing in terms of another" (Semino, 2008, p. 1). Conceptual Metaphor Theory foregrounded metaphor as a mechanism of thought that structures everyday conceptualisation (Lakoff & Johnson, 2003). Metaphors are thus pervasive conceptual tools capable of structuring and restructuring reality, which makes them highly valuable for fostering critical AI literacy and explaining opaque generative AI systems (Roe et al., 2025). Subsequent work has elaborated this position by emphasising that metaphor is simultaneously cognitive, linguistic, and communicative: what counts as metaphor in use depends on discourse context, genre, and audience (Steen, 2008, 2011; Deignan et al., 2013).

A recurring critique of early conceptual accounts of metaphor concerns over-reliance on introspection and insufficient attention to authentic and representative linguistic evidence to substantiate the posited conceptual metaphors. This has driven the development of more data-driven approaches and, crucially, systematic identification methods. MIP and MIPVU provide operational criteria for determining metaphorically used words in context by contrasting contextual meaning with a more basic meaning (Pragglejaz Group, 2007; Steen et al., 2010). These procedures have enabled more transparent metaphor research, especially in corpus and discourse studies, and they also provide a robust analytical anchor for elicited data, where the unit of analysis is often a short response rather than an extended discourse segment.

Metaphor elicitation is an established technique for capturing speakers' preferred conceptualisations in a comparatively controlled manner (Wan & Low, 2015). Standard prompts such as "X is like Y" support systematic comparison, but elicitation also introduces familiar methodological challenges, including how to treat similes, how to avoid conflating metaphor with metonymy, and how to account for task framing effects (Wan & Low, 2015; Low, 2015). For these reasons, elicitation studies typically combine careful prompt design with explicit analytic procedures, often integrating qualitative interpretation with quantitative distributional summaries (Low, 2015; Strugielska, 2015).

A targeted review of brain metaphors shows continuity, but also inevitable shifts across historical periods. Earlier scientific and philosophical discourse frequently relied on mechanical and spatial analogies (Crane, 2003; Dupuy, 2009), while late twentieth-century cognitive science normalised computational framings of mind and brain (Draaisma, 2000; Hayles, 1999). Computational metaphors have been productive, but their dominance has also been questioned, especially where they risk being mistaken for literal descriptions rather than modelling choices (Brette, 2019, 2022). Prior work also highlights recurring metaphor families such as (i) reification and spatialisation (e.g., memory as storage), (ii) personification (attributing agency to the brain), and (iii) technological metaphors (e.g., the brain as an engineered system) (Draaisma, 2000; Goschler, 2007). These patterns provide a comparative background for analysing what humans and LLMs produce when asked to generate brain metaphors.

Finally, the question of metaphorical competence in AI has increasingly become an empirical topic. In applied and computational work, metaphor processing is often divided into identification, interpretation, and generation (Ge et al., 2023). Recent studies suggest that LLMs can interpret at least some novel metaphors in constrained tasks (Ichien et al., 2024), while comparative work indicates differences between human and LLM metaphor production, including a tendency towards more conventional sources or "closer" mappings (Mao et al., 2024). These studies motivate direct, task-matched comparisons using established metaphor-analytic tools. The present study contributes by applying MIPVU-based analysis to elicited brain metaphors from humans and multiple LLM-based systems, thereby situating model outputs within a framework commonly used in metaphor studies.

Brain Metaphors in Scientific Discourse: A Brief Review of Previous Scholarship

Scholars have documented an evolving array of metaphors for the brain in scientific and popular discourse. Early brain research drew heavily on mechanical and spatial analogies – brains as pumps, telegraphs, or switchboards – reflecting the technological milieu of the time (Crane, 2003; Dupuy, 2009). With the rise of computing, digital metaphors became predominant, describing the brain as hardware or software, a serial processor or a neural network. These metaphors illuminated some aspects of neural function while obscuring others, prompting debates on their epistemic merits and pitfalls (Brette, 2019). Just as the brain has historically been compared to machines, modern AI algorithms are reciprocally described using anthropomorphised metaphors, adopting human cognitive terms like "learning," "understanding," and "deciding" (Rehak, 2021). While these metaphors make the technology approachable, they also perpetuate deterministic myths, obscure the mathematical reality of these systems, and distance the output from human responsibility (Garibay-Petersen et al., 2025; Roe et al., 2025).

Reification/Spatialisation, Personification, and Technological Metaphors

Existing reviews highlight three broad families of brain metaphors (Draaisma, 2000; Goschler, 2007). Reification/spatialisation metaphors treat abstract cognitive processes as concrete objects or spaces (e.g., memory as a storage container, knowledge as a substance in a vessel). Personification metaphors attribute human-like agency or characteristics to the brain ("the brain remembers," "the brain decides"), implicitly separating the organ's function from the person's self and sometimes fostering a dualistic outlook. Technological metaphors, from clocks to computers, have been especially influential, embedding assumptions of modularity, control, and predictability. Each category frames scientific questions differently: spatial metaphors suggest locating functions, personification raises questions of autonomy, and technological metaphors bias thinking toward design and engineering concepts. To counter uncritical personification and technological assumptions, recent educational literature advocates for deliberate, critical metaphors to describe AI, such as the “black box” (highlighting algorithmic opacity), the “echo chamber” (emphasising the reinforcement of biases), or the “funhouse mirror” (illustrating how AI distorts reality) (Roe et al., 2025).

Expanding Metaphor Families beyond Computation

Recent scholarship has begun cataloguing less dominant metaphors that push beyond the computational paradigm (Brette, 2022; Matassi & Martinez, 2023). These include organic and ecological metaphors (the brain as a forest, an ecosystem, a garden) emphasising growth, adaptation, and interconnectedness; social metaphors (the brain as a society or city) highlighting coordination among semi-autonomous units; and artisanal metaphors (the brain as a workshop or library) which underscore knowledge accumulation and skill. While rare in mainstream literature, such metaphors exist in niche discussions and may carry pedagogical advantages by connecting neuroscience concepts to embodied, familiar experiences. Their limited uptake suggests that dominant technological narratives continue to overshadow alternative framings. Furthermore, scholars argue that emerging technologies require entirely new conceptual metaphors; for instance, GenAI has been conceptualised as “digital plastic”, i.e., a ubiquitous, synthetic, and highly malleable resource that democratises creation but risks polluting the digital ecosystem with homogenous, culturally biased content if not critically examined (Roe et al., 2025).

Historical and Cultural Factors Affecting Contemporary Usage

The prominence of certain brain metaphors is partly historically contingent. For example, the dominance of computer metaphors in the late 20th century coincided with the digital revolution and cognitive science’s information-processing paradigm (Dupuy, 2009). Culturally, different languages and traditions emphasise different analogies; therefore, recognising these influences is crucial for a critical understanding of how metaphors shape and sometimes limit scientific imagination. It also underscores the importance of introducing a plurality of metaphors, particularly in educational contexts, to avoid calcifying one model as the only way to think about the brain. It must also be acknowledged that GenAI models themselves act as “culture machines” that lean towards Western-centric and Eurocentric linguistic and cultural generalities embedded in their training data (Jones, 2025; Roe, 2025). Consequently, metaphors generated or interpreted by AI may lack accessibility or relevance for users from different cultural backgrounds (Roe et al., 2025). Furthermore, educators and researchers must recognise that metaphors possess conceptual limits and can “break down” when extended too far, meaning their cultural appropriacy and limitations must be continually evaluated (Carter & Pitcher, 2010; Roe et al., 2025).

Metaphorical Competence of Artificial Intelligence

Metaphor scholars define metaphorical competence as the knowledge and capacity to use and interpret metaphors (Littlemore & Low, 2006, p. 269). Littlemore (2001, p. 461) identifies four key elements in metaphor competence, namely, (a) ability to create original metaphors, (b) interpretive fluency, (c) ability to extract their meaning, and (d) the speed with which that meaning is identified. While humans typically manifest these dimensions through embodied cognition, the extent to which AI approximates them remains uncertain. Indeed, the very notion that generative AI can be “creative” in a human-like manner is hotly contested, reflecting scepticism among human creative practitioners about whether AI-generated language truly exhibits original imagination. Recent research suggests that AI language models can identify and interpret conventional metaphors but struggle with creative generation. Ichien et al. (2024, p. 296) caution that while AI can interpret novel metaphors, this does not equate to the capacity to invent them – its production appears confined to recombination of human-authored patterns. These findings highlight key deficits in AI’s metaphorical competence. Mao et al. (2024) expand on this by providing a comparative cognitive lens: using MetaPro, they contrast metaphor usage between ChatGPT and humans, showing that ChatGPT favours more conventional domains and less ‘distant’ mappings, indicating potential constraints in its conceptual flexibility. It lacks novelty, embodiment, and emotional resonance, which can be considered the cornerstones of human metaphorical creativity. Similarly, outside the metaphor domain, ChatGPT-4’s outputs have been shown to be highly polished yet less diverse than human creations – for example, its advertising slogans were rated more novel and persuasive than those by humans, but exhibited significantly lower variety (Gill & Hoxha, 2024). These deficits are compounded by the fact that AI models lack true phenomenological experience or "qualia," offering instead a textuality rooted strictly in symbolism and statistical probability rather than lived experience (van Manen, 2023; Roe, 2025). This limitation is famously captured by the "stochastic parrot" metaphor, which argues that AI merely reassembles content from its training data based on probabilities without any authentic understanding (Bender et al., 2021; Roe et al., 2025). If AI-generated outputs are continually recycled into future training data without true conceptual grounding, it acts as "digital microplastics" that can lead to "model collapse" and degrade the quality of meaning-making ecosystems (Shumailov et al., 2024; Roe et al., 2025).

Data and Methodology

This study adopts an elicited-metaphor design to compare how humans and LLM-based systems metaphorically describe the brain and what conceptual patterns each set evoke. Data from both groups were collected using closely matched prompts, enabling comparison between two sets of metaphor producers.

Human Participants and Sampling

Human metaphor data were gathered via a Google Forms questionnaire distributed through snowball sampling (Lewin, 2005, p. 219). The sample comprised 61 participants. Demographically, the group was predominantly aged 18–27 (approx. 70%) and predominantly female (approx. 70%). Most participants were L2 users of English. The most frequent L1 backgrounds were Lithuanian (30%), Russian (25%), Dutch (20%), and English (16%). This demographic profile is relevant when considering external validity: metaphor choice and linguistic realisation can be shaped by language background, educational experience, and familiarity with scientific discourse (Semino, 2008; Low, 2008).

Questionnaire Prompts and Procedure

The questionnaire employed open-ended sentence stems designed to elicit metaphorical comparisons of the brain. Participants responded to the following five prompts:

The brain is…

The brain is like…

The brain can be compared to…

The brain functions like…

The brain has functions similar to…

The term "metaphor" was avoided in the instructions. Instead, participants were guided using more accessible terms such as "comparison" or "analogy", reflecting findings that non-specialist respondents may associate "metaphor" primarily with literary or poetic language (Low, 2008). To capture figurative construals that might not be overtly marked as comparisons, participants also provided a short description of how the brain works "without using analogies," following elicitation strategies used to surface implicit metaphoricity (Cameron & Maslen, 2010). Participants were instructed not to consult generative AI systems or other external sources while completing the task and were asked to self-confirm compliance. Responses were collected anonymously.

AI Metaphor Generation

To enable systematic comparison, metaphors were elicited from eight widely used LLM-based chat platforms: ChatGPT, Gemini, Copilot, Claude, Copy, Chatsonic, Perplexity, and Pi. The same prompt set used for the human questionnaire was used for each platform. Each platform was queried five times, with sessions spaced over four days. Repeated querying was used to reduce the likelihood that results reflected a single session's output idiosyncrasies and to sample variability across outputs, noting that LLM responses can differ across time and session context (Chang & Bergen, 2024).

To limit contextual carry-over, queries were run in minimal-context sessions (i.e., without extended prior conversation). The recorded output for each query consisted of the platform's primary metaphorical comparison and any accompanying explanation.

Analytical Procedure

The analytical unit was a metaphorical description of the brain (e.g., "The brain is like a city"), treated as a single item that evoked a coherent metaphor within a classification. From the human dataset, 66 metaphor items were extracted from the full set of questionnaire responses (some responses contained multiple comparisons). The AI dataset comprised 40 items (8 platforms x 5 sessions). Where AI response used additional nested comparisons to elaborate a main cross-domain mapping (e.g., "brain as city" plus correspondences such as "hippocampus as library"), the main source domain introduced by the response was used for quantitative categorisation; nested correspondences were retained for qualitative analysis of explication strategies.

All responses were analysed applying the main principles of the Metaphor Identification Procedure Vrije Universiteit (MIPVU) (Steen et al., 2010). In brief, potentially metaphor-related words were examined by comparing their contextual meaning with a more basic meaning (if relevant) and determining whether the contextual meaning could be understood via cross-domain mapping. In the first round, both analysts identified metaphorical items individually, and any discrepancies were resolved in subsequent rounds of discussion.

Following the identification, each item identified as evoking a brain metaphor was coded for the source domain(s) that the linguistic items evoke. Thirteen source-domain categories were used, adapted from prior discussions of brain metaphors and refined through inductive pilot coding (e.g. Draaisma, 2000; Goschler, 2007; Brette, 2019). The final taxonomy was established through iterative double coding by two independent raters, with discrepancies resolved through discussion and consensus, ensuring methodological transparency.

Results and Discussion

This section integrates quantitative and qualitative findings to examine how humans and large language models (LLMs) conceptualise the brain through metaphor. Beyond frequency counts, the analysis explores how metaphor choice reflects underlying differences in cognitive grounding, elaborative strategy, and metaphorical competence.

Quantitative Overview of Source Domain Distribution

Across the combined dataset of 106 metaphors (66 human-derived, 40 AI-derived), a broad trend was evident: both humans and AIs relied heavily on technological source domains, but with divergent secondary preferences. Technology-related metaphors accounted for 46.97% of all human metaphors and 35% of AI metaphors. However, based on the functions highlighted in the qualitative explanations, all were assigned to different higher-level semantic categories. As shown in Figure 1, Network-related metaphors comprised more than 45.00% of AI outputs but only 9.09% of human outputs. Moreover, the Information processor category was notable among both human participants and LLMs, accounting for 16.67% and 30% of generated source domains, respectively. Management metaphors were largely present in both (15.15% human; 25% in AI, where many "supercomputer" and "control centre" responses combined elements of tech and management). There were also a number of semantic domains, representing less than 5% of source domains, suggested by human participants (natural phenomena, building, container, etc.). They were also virtually absent in AI metaphors (only 2 instances, or 5%, fell outside technology or network domains). Table 1 summarises the principal metaphor families identified in both datasets, including the primary aspects by which the source domains were assigned to them. Five source domain categories, i.e. Mechanism/Device, Information repository, Information processor, Management, and Network, accounted for most instances across human and AI outputs.

Table 1 Comparison of source domains used in human- vs AI-generated metaphors

Metaphorical category	Human Count (out of 66)	AI Count (out of 40)	Source Domains Used by Humans	Source Domains Used by AI	Category explanations
Management	10	10	CONTROL CENTRE, CPU, CONTROLLER / COMMANDER, COMPUTER, MOTHERBOARD, ROBOT, SYSTEM UNIT, FACTORY EMPLOYEE	ORCHESTRA CONDUCTOR, SYMPHONY ORCHESTRA, ORCHESTRA, CONTROL CENTRE, MOTHERBOARD, CITY, ORCHESTRA, CITY'S CONTROL CENTRE	An entity responsible for coordinating, directing, or controlling other components ➔ used when the brain is described as being responsible, "in charge," or organising activities *Focus: leadership, control, and regulation.*
Information repository	12	0	FILE CABINET, BOOK, STORAGE UNIT, COMPUTER, HARD DRIVE, BLACK BOX, LIBRARY, HARD DISK, DATA CENTRE		A system which main purpose is to store and retain information over time ➔ used when the analogy highlights memory, storage, or accumulation of knowledge *Focus: memory, storage, and keeping information over time.*
Mechanism / device	11	0	CLOCK, PROGRAMMED MECHANISM, MECHANISM / MACHINE, NUCLEAR REACTOR, COMPUTER, QUANTUM COMPUTER		A human-made object or machine designed to perform a function though physical or technical operation ➔ used when the analogy emphasises how something works mechanically or technically *Focus: how something works like a machine or device.*
Network	6	18	NETWORK, CITY, LABYRINTH, COMPLEX SYSTEM, KNITTING YARN,	CITY, HIGHWAY, NETWORK OF COMPUTERS, NETWORK OF ROADS AND HIGHWAYS	A system defined by interconnected elements and the transfer of signals or information between them ➔ used when the analogy emphasised connections, pathways, or communication flow *Focus: connections, pathways, communication, and signal transmission.*
Information processor	11	12	COMPUTER, AI, SOFTWARE, OPERATING SYSTEM CONTROL SYSTEM	COMPUTER, SUPERCOMPUTER, CPU	A system whose main purpose is to interpret, transform, or compute information ➔ used when the analogy focuses on thinking, calculating, or decision-making processes *Focus: thinking, decision-making, computation, and analysis.*

Metaphorical category	Human Count (out of 66)	AI Count (out of 40)	Source Domains Used by Humans	Source Domains Used by AI	Category explanations
Human body parts	3	0	MUSCLE		Analogies comparing brain to specific physical parts of the body ➜ used when the comparison is to a distinct organ or body structure, not a whole organism *Focus: a single organ (like the heart, brain, or eye) rather than the whole body.*
Household items	1	0	KETTLE		Everyday objects found in domestic settings used for practical purposes. ➜ used when the analogy relies of familiar home-use objects, not their technical functions *Focus: simple, familiar items (not technical in function).*
Container	3	0	SPONGE, POT		An object whose defining feature is that it holds or contains something within it. ➜ used when the analogy emphasises capacity, enclosure, or holding, not storing information specifically. *Focus: capacity, containment, and boundaries.*
Natural phenomena	3	0	UNIVERSE, BODY OF WATER, GALAXY		Processes or systems that occur naturally without human design. ➜ used when the analogy focuses on organic, or environmental processes. *Focus: natural processes like growth, weather, or evolution.*
Individual organism	2	0	PERSON		A complete living being, considered as a whole system. ➜ used se when the brain is compared to an entire organism, not just a body part. *Focus: the entire organism and its life processes.*
Building	2	0	OFFICE BUILDING, FACTORY		A large, constructed structure with defined spaces, designed for people to occupy or use. ➜ use when the analogy emphasises rooms, sections, or internal layout, organisation of space, or different areas serving different functions *Focus: layout, rooms, and spatial structure.*

Metaphorical category	Human Count (out of 66)	AI Count (out of 40)	Source Domains Used by Humans	Source Domains Used by AI	Category explanations
Food	1	0	SOUP		Any edible substance or cooking-related item/process. ➜ used when the analogy relates to ingredients or any type of food. *Focus: ingredients, edible substances.*
Agriculture	1	0	PILE OF HAY		Systems related to cultivation, farming, or growth over time. ➜ used when the analogy employs any agricultural domains. *Focus: agricultural products, farming practices, and growth-related contexts.*

Overall, Table 1 reveals distinct tendencies in how humans and AI systems conceptualise the brain through metaphor. Human participants produced a more diverse and experientially grounded range of source domains, extending beyond technological and organisational framings to include embodied, domestic, and natural imagery. By contrast, AI-generated metaphors were more systematic and abstract, concentrating heavily on network and management domains that mirror the architectural logic of computational and informational systems. Although both groups drew on the entrenched brain as computer metaphor family, their elaborations diverged: human responses distributed this model across multiple subdomains – mechanical, managerial, and mnemonic – while LLMs' outputs condensed around coherent system-level analogies, realising such source domains as CITY, ORCHESTRA, and SUPERCOMPUTER. In sum, the findings suggest that human metaphorisation remains largely grounded in lived experience and material analogies, while AI's figurative output reflects structural reasoning rather than embodied understanding, indicating complementary but asymmetrical patterns of metaphorical competence.

Analysis of Brain Metaphors Used by Humans versus AI: Quantitative Tendencies

The comparative distribution of categories revealed clear differences in metaphorical preference and scope. All responses were classified into 13 higher-level semantic categories based on the aspects of analogy highlighted in the qualitative explanation (see **Table 1** and **Fig. 1**). Additionally, some source domains generated were further grouped under a broader Technological source domain grouping to capture the frequency of technological domains in brain metaphorical conceptualisations, despite variation in the specific functions emphasised.

Human responses cluster around technology-related representations (46.97% of all human metaphors). However, despite being combined by a shared technological source domain grouping, these conceptualisations emphasise different brain functions. As can be observed from Figure 1, the category of the Information repository dominates human responses (18.18%), highlighting the brain's ability to store and retain information over time. This is closely followed by Information processor, which focuses on thinking, decision making and analysis, and Mechanism/device, which emphasises general technological functioning; both categories account for 16.67%. The last prominent category is Management (15.15%), describing the brain as responsible for coordinating and organising other parts of the body.

By contrast, technological metaphors account for only 35% of AI outputs. Instead, AI predominantly describes the brain through network metaphors which constitute 45% of all outputs. Across all source domain categories, the Network category is the most frequent (45%), realised through metaphors of the brain as a CITY, HIGHWAY or another infrastructural source domains that highlight connectivity.

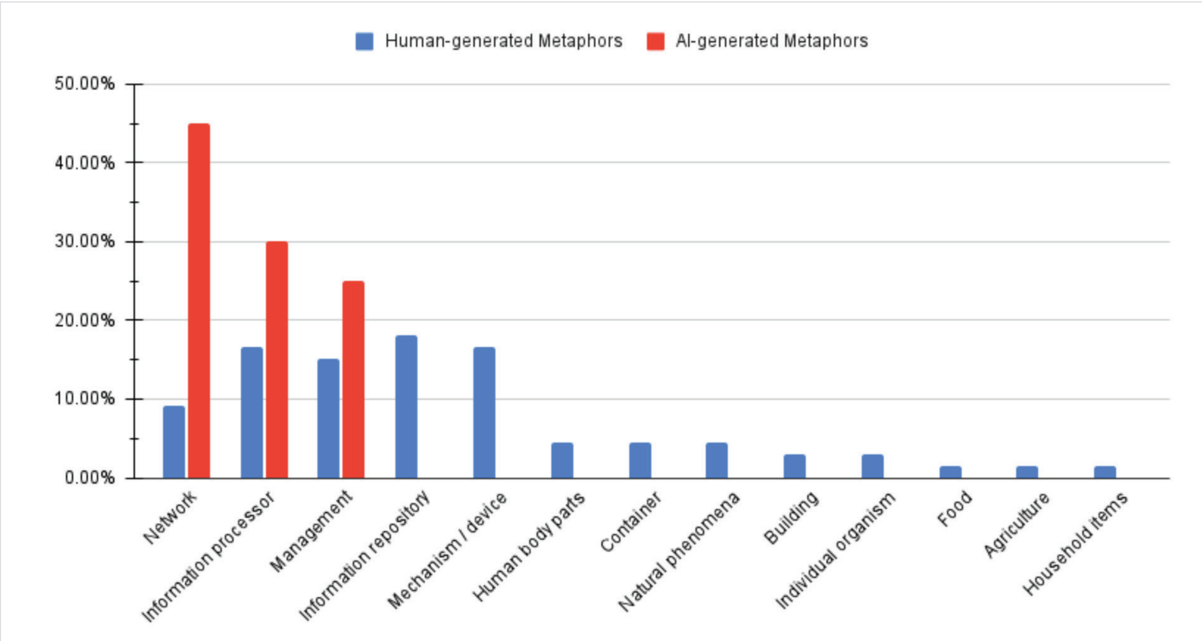

Fig. 1 Distribution of the source domain groupings in human vs AI-generated metaphors (%)

Viewed comparatively, the human set exhibits greater semantic breadth. In addition to technological and organisational metaphors, participants employed low-frequency but experientially grounded source domains, including human body parts, household items, containers, natural phenomena, buildings, individual organisms, food, and agriculture. Although each of these categories accounts for less than 5% of the data, collectively they are indicative of a wide exploratory search of source domains derived from everyday experience to describe the brain. The AI set shows semantic concentration, with negligible representation beyond network and technology, and no instances in the embodied or domestic source domains.

Human-produced Metaphors

Within the human dataset, Information repository metaphors (18.18%) constitute the dominant category, typified by FILE CABINET, BOOK, and STORAGE UNIT. These analogies foreground the brain as a storage and retrieval system, suggesting that many participants conceptualise the brain primarily as a digital repository. Network metaphors (9.09%), exemplified by NETWORK, CITY, and LABYRINTH, depict the brain as a structure composed of densely interconnected elements, emphasising communication among components as a condition for integrated functioning.

Less frequent categories (each below 5%) enrich this conceptual repertoire, including human body parts (BRAIN IS A MUSCLE), household items (BRAIN IS A KETTLE, BRAIN IS A SPONGE), natural phenomena (BRAIN IS A BODY OF WATER), food (BRAIN IS SOUP), agriculture (BRAIN IS A PILE OF HAY), and buildings (OFFICE BUILDING). For example, building analogies depict the brain as a physical structure with defined areas responsible for different functions. Though it was infrequent, participants sometimes elaborated such comparisons through extended analogies that linked organisational structure with neural functioning. For example, one participant described the brain as “a massive unified office building with nearly endless number of working departments, warehouses and garages with company cars,” explaining: “I just tried to find comparisons and remembered my job. Then, when you think about it, departments would be different parts of the brain that each respond to different things, warehouses are where information is stored in the brain, and cars are like impulses (I don’t know anything about biology) that communicate with the body”. This explication demonstrates functional reasoning emerging from everyday experience: the metaphor not only maps physical structure onto mental architecture but also interprets neural communication through the explicit comparison of workplace logistics. Such elaborations reveal how participants actively integrate personal and occupational semantic components into metaphor construction. Although these occur sporadically, they extend the representational range of the dataset, reflecting an embodied and domestic orientation to cognition.

Furthermore, the examination of the metaphors which can be united based on their technological aspects, represented in **Fig. 2**, shows a more differentiated pattern. Four recurrent groups from the prior defined semantic families emerged: Mechanism/Device (32.3%), Information processor (32.3%), and smaller Management (19.4%) and Information repository (16.1%) subgroups, as indicated in **Fig. 2**. Human participants most frequently employed COMPUTERS and PARTS OF A COMPUTER (CPU, motherboard, hard drive) as source domains for the brain.

The graphical distribution highlights that human-produced technological metaphors balance mechanical and managerial categories rather than being dominated by a single pattern. Respondents often described the brain as a mechanical or organisational system, yet they also introduced analogies related to data storage and recall, capturing the interplay between processing and memory. In effect, **Fig. 2** visualises a cognitive pluralism in human metaphorisation: a roughly even division between mechanism and management, complemented by smaller but conceptually significant repositories and processors.

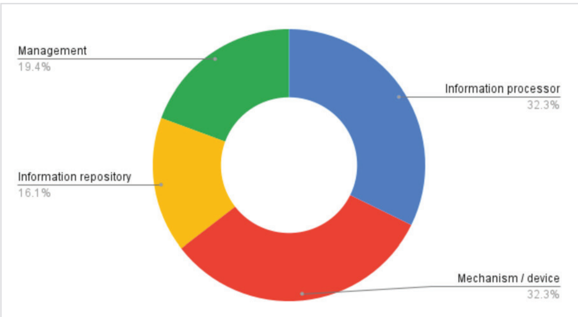

Fig. 2 Distribution of the source domain subgrouping in human-generated metaphors in Technology category (%)

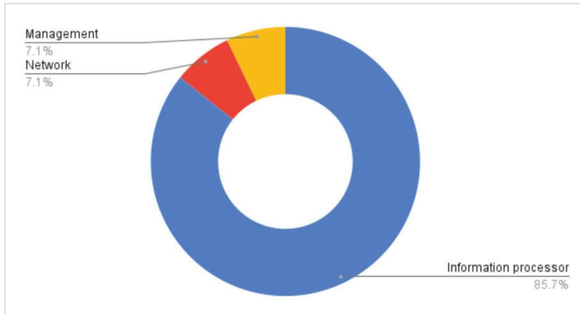

Fig. 3 Distribution of the source domain subgrouping in AI-generated metaphors in Technology category (%)

Technological metaphors, while being the largest group, were far from uniform. Many participants chose the ubiquitous "brain is a computer" analogy, but others specified variations (the brain as a calculator for its precision, as a database for memory storage, or as a programmer when emphasising control). Notably, even when using similar base comparisons, human respondents tended to add idiosyncratic twists or examples from personal experience. One described the brain as "like a smartphone that has many apps running – sometimes you need to close some to focus on one task." Here the participant extended the computer metaphor into a contemporary device and implicitly addressed cognitive load and multitasking.

AI-generated Metaphors

Network metaphors were the most prominent in the AI set, accounting for 45.00% of all AI-generated source domains. These conceptualise the brain as a CITY, HIGHWAY, or NETWORK OF COMPUTERS, foregrounding coordination, interaction, and the interconnectedness of parts. This emphasis aligns with the biological reality of millions of interconnected neurons forming large-scale circuits. The Information processor set constitutes the second largest group (30.00%), including source domains such as COMPUTER, SUPERCOMPUTER, and CPU, which emphasise the brain's role in the interpretation and transformation of information, as well as the further execution of commands.

Domains employing technological conceptualisations constitute 35% of all cases, consisting primarily of the Information processor category (e.g., BRAIN IS A COMPUTER), and including instances from Management (e.g., BRAIN IS A MOTHERBOARD) and Network (e.g., NETWORK OF COMPUTERS) categories at a higher semantic level. As **Fig. 3** shows, these analogies highlight accuracy, efficiency, and processing/control capabilities, reflecting a tendency to map neural function onto engineered systems. In line with the human data, the most frequent concrete sources within this family were COMPUTER and SUPERCOMPUTER, with smaller contributions from PART OF A COMPUTER and NETWORK OF COMPUTERS.

A further Management strand (25%) – often expressed through culturally salient organisers such as an ORCHESTRA / ORCHESTRA CONDUCTOR or A CITY'S CONTROL CENTRE – again represents the brain as a central coordinating entity that harmonises many parts toward common goals. Overall, the AI profile is thus systemic and organisational: it privileges networked infrastructure and engineered control as default frames and, characteristically, accompanies them with explicit explanatory mapping.

Conclusions

This study aimed to provide an empirically grounded comparison of how humans and LLMs describe the brain through metaphor. By using a matched elicitation design, the research addressed three main questions regarding most recurrent source domains evoked, diversity of the domains across the two samples and the extent to which humans and AI explicitly articulate the proposed cross-domain mappings.

Regarding the most recurrent source domains evoked, both human participants and LLMs relied heavily on technological metaphors, especially the deeply entrenched and conventional BRAIN IS A COMPUTER metaphor. However, their secondary preferences diverged rather significantly. Human responses distributed the technological source domains across a balanced range of sub-domains, predominantly favouring the Information repository (e.g. FILE CABINET, STORAGE UNIT), Information processor, Mechanism/device, and Management groupings. By contrast, LLM-generated outputs concentrated heavily on Network metaphors (e.g. CITY, HIGHWAY, NETWORK OF COMPUTERS, etc.) which accounted for nearly half of the AI responses, followed by Information processor and Management groupings.

When comparing the datasets in terms of their source-domain ranges and diversity, the findings reveal a distinct asymmetry. Human metaphorisation demonstrated a significantly broader and more diverse semantic range of the source domains for the conceptualisation of the brain. Alongside technology-based descriptions, human participants frequently drew on idiosyncratic, embodied, and experientially grounded source domains, incorporating references to human body parts (a muscle), household items (a sponge, a kettle), natural phenomena, and agriculture. Conversely, AI-generated metaphors were much more homogenous and abstract, showing an almost complete absence of embodied or domestic categories. Instead, AI systems consistently preferred coherent, system-level organisational analogies that reflect structural reasoning rather than human lived experience.

In terms of explicitness and elaboration, the two sets exhibited diverging communicative orientations. LLMs articulated cross-domain mappings with a high degree of explicit justification and elaboration, producing responses that were approximately five times longer on average than those of human participants. The AI outputs demonstrate outstanding linguistic fluency and systematic structural mapping. Human participants, however, favoured interpretive openness and concision, relying more on tacit, affective, and contextually rich associations rather than exhaustive metalinguistic explication prevalent in LLM-based outputs.

Synthesising these findings, the comparative analysis highlights complementary strengths of the two samples. Current LLMs seem to function as rather accurate expositors of conventional, structurally coherent analogies, yet they simultaneously exhibit limited metaphorical competence regarding source domain variation, metaphor novelty and embodied grounding. Human metaphorical competence seems to remain largely rooted in physical, embodied and other experiential domains, while AI outputs primarily recombine and systematically articulate the architectural and infrastructural patterns prevalent in their training data.

Several limitations to this study must be acknowledged. First, the human dataset was derived from a relatively small and demographically specific sample, consisting predominantly of young, female L2 English speakers from a narrow range of L1 backgrounds. This demographic profile may have influenced the specific linguistic realisations and cultural metaphor choices observed. Second, AI data reflect a temporal snapshot of LLM capabilities across various proprietary architectures as of 2024-2025; given the rapid evolution of generative AI, these outputs are subject to change. Third, despite the use of multiple rounds of double-coding procedures, the evaluation of metaphor's creativity, embodiment, and conceptual aptness retains an inherent degree of interpretive subjectivity.

Future research should seek to expand the human dataset across broader age groups, disciplines, and linguistic backgrounds to enhance cultural generalisability. Introducing experimental manipulations such as explicitly prompting LLMs to generate embodied or novel metaphors could help determine whether AI's reliance on structural mapping is an inherent cognitive constraint or a reversible output bias. Finally, integrating multimodal and cross-linguistic elicitation tasks would provide a more comprehensive understanding of human versus artificial figurative reasoning.

Conflict of Interest

The authors declare no conflict of interest regarding the publication of this article.

Declaration on the Use of AI

During the preparation of this chapter, the authors used generative AI tools solely as part of the research methodology for the elicitation of metaphors. The application of AI in this process is fully described in the Data and Methodology section. After using the tools, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

1 Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021). On the dangers of stochastic parrots: Can language models be too big? Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21), 610-623. https://doi.org/10.1145/3442188.3445922

2 Brette, R. (2019). Is coding a relevant metaphor for the brain? Behavioral and Brain Sciences, 42, e215: 1-58. https://doi.org/10.1017/S0140525X19000049

3 Brette, R. (2022). Brains as computers: Metaphor, analogy, theory or fact? Frontiers in Ecology and Evolution, 10, Article 878729. https://doi.org/10.3389/fevo.2022.878729

4 Cameron, L., & Maslen, R. (Eds.). (2010). Metaphor Analysis: Research Practice in Applied Linguistics, Social Sciences and the Humanities. London: Equinox.

5 Carter, S., & Pitcher, R. (2010). Extended metaphors for pedagogy: using sameness and difference. Teaching in Higher Education, 15(5), 579-589. https://doi.org/10.1080/13562517.2010.491904

6 Chang, T. A., & Bergen, B. K. (2024). Language model behavior: A comprehensive survey. Computational Linguistics, 50(1), 293-350. Retrieved September 23, 2025, from https://aclanthology.org/2024.cl-1.9.pdf. https://doi.org/10.1162/coli_a_00492

7 Crane, T. (2003). The Mechanical Mind: A Philosophical Introduction to Minds, Machines and Mental Representation (2nd ed.). Routledge.

8 Deignan, A., Littlemore, J., & Semino, E. (2013). Figurative Language, Genre and Register. Cambridge University Press.

9 Draaisma, D. (2000). Metaphors of memory: A history of ideas about the mind. Cambridge University Press.

10 Dupuy, J.-P. (2009). The mechanization of the mind: On the origins of cognitive science. Princeton University Press.

11 Garibay-Petersen, C., Lorimer, M., & Menzat, B. (2025). Creating certainty where there is none: Artificial intelligence as political concept. Big Data & Society, 12(4). https://doi.org/10.1177/20539517251396079

12 Ge, M., Mao, R., & Cambria, E. (2023). A survey on computational metaphor processing techniques: From identification, interpretation, generation to application. Artificial Intelligence Review, 56, 1829-1895. https://doi.org/10.21203/rs.3.rs-2220808/v1

13 Gill, T., & Hoxha, K. (2024). Chat-GPT4 can enhance marketing creativity. But human ego can hamper its potential. [Preprint]. SSRN. https://doi.org/10.2139/ssrn.5135496

14 Goschler, J. (2007). Metaphors in cognitive and neurosciences: Which impact have metaphors on scientific theories and models? metaphorik.de, 12, 63-80. https://doi.org/10.17185/duepublico/83554

15 Hayles, N. K. (1999). How we became posthuman: Virtual bodies in cybernetics, literature, and informatics. University of Chicago Press. https://doi.org/10.7208/chicago/9780226321394.001.0001

16 Ichien, N., Stamenković, D., & Holyoak, K. J. (2024). Large language model displays emergent ability to interpret novel literary metaphors. Metaphor and Symbol, 39(4), 296-309. https://doi.org/10.1080/10926488.2024.2380348

17 Jones, R. H. (2025). Culture machines. Applied Linguistics Review, 16(2), 753-762. https://doi.org/10.1515/applirev-2024-0188

18 Lakoff, G., & Johnson, M. (2003). Metaphors we live by. University of Chicago Press. https://doi.org/10.7208/chicago/9780226470993.001.0001

19 Lewin, C. (2005). Elementary quantitative methods. In B. Somekh & C. Lewin (Eds.), Research methods in the social sciences (pp. 215-225). Sage.

20 Littlemore, J. (2001). Metaphoric competence: A language learning strength of students with a holistic cognitive style? TESOL Quarterly, 35(3), 459-491. https://doi.org/10.2307/3588031

21 Littlemore, J., & Low, G. (2006). Figurative thinking and foreign language learning. Palgrave Macmillan. https://doi.org/10.1057/9780230627567

22 Low, G. (2008). Metaphor and education. In R. W. Gibbs, Jr. (Ed.), The Cambridge handbook of metaphor and thought (pp. 212-231). Cambridge University Press. https://doi.org/10.1017/CBO9780511816802.014

23 Low, G. (2015). A practical validation model for researching elicited metaphor. In W. Wan & G. Low (Eds.), Elicited metaphor analysis in educational discourse (pp. 15-38). John Benjamins Publishing Company. https://doi.org/10.1075/milcc.3

24 Matassi, G., & Martinez, P. (2023). The brain-computer analogy - "A special issue." Frontiers in Ecology and Evolution, 10. https://doi.org/10.3389/fevo.2022.1099253

25 Mao, R., Chen, G., Li, X., Ge, M., & Cambria, E. (2024). A comparative analysis of metaphorical cognition in ChatGPT and human minds. Cognitive computation. https://doi.org/10.1007/s12559-024-10393-y

26 Muñoz-Ortiz, A., Gómez-Rodríguez, C., & Vilares, D. (2024). Contrasting linguistic patterns in human and LLM-generated news text. Artificial Intelligence Review, 57, Article 265. https://doi.org/10.1007/s10462-024-10903-2

27 Pragglejaz Group. (2007). MIP: A method for identifying metaphorically used words in discourse. Metaphor & Symbol, 22(1), 1-39. https://doi.org/10.1080/10926480709336752

28 Rehak, R. (2021). "Action" and ascription: On misleading metaphors in the debate about artificial intelligence and transhumanism. In H. T. Engelhardt Jr. & B. Gordijn (Eds.), Transhumanism: The proper guide to a posthuman condition or a dangerous idea? (pp. 191-210). Springer. https://doi.org/10.1007/978-3-030-56546-6_11

29 Roe, J. (2025). How to use generative AI in educational research. Cambridge University Press. https://doi.org/10.1017/9781009675338

30 Roe, J., Furze, L., & Perkins, M. (2024). Reflecting reality, amplifying bias? Using metaphors to teach critical AI literacy. Journal of Interactive Media in Education, 1, 1-15. https://doi.org/10.5334/jime.961

31 Semino, E. (2008). Metaphor in discourse. Cambridge University Press. https://doi.org/10.1017/CBO9780511816802.015

32 Shumailov, I., Shumaylov, Z., Zhao, Y., Papernot, N., Anderson, R., & Gal, Y. (2024). AI models collapse when trained on recursively generated data. Nature, 631, 755-759. https://doi.org/10.1038/s41586-024-07566-y

33 Steen, G. (2008). The paradox of metaphor: Why we need a three-dimensional model of metaphor. Metaphor and Symbol, 23(4), 213-241. https://doi.org/10.1080/10926480802426753

34 Steen, G., Dorst, A., Herrmann, J. B., Kaal, A., Krennmayr, T., & Pasma, T. (2010). A method for linguistic metaphor identification: From MIP to MIPVU. John Benjamins. https://doi.org/10.1075/celcr.14

35 Steen, G. (2011). The contemporary theory of metaphor - Now new and improved! In F. González-García, M. S. Peña Cervel, & L. Pérez Hernández (Eds.), Metaphor and metonymy revisited beyond the contemporary theory of metaphor (pp. 27-66). John Benjamins. https://doi.org/10.1075/rcl.9.1.03ste

36 Strugielska, A. (2015). A hybrid methodology of linguistic metaphor identification in elicited data and its conceptual implications. In W. Wan & G. Low (Eds.), Elicited metaphor analysis in educational discourse (pp. 65-89). John Benjamins Publishing Company. https://doi.org/10.1075/milcc.3.03str

37 Umair, M., Mertens, J. B., Warnke, L., & de Ruiter, J. P. (2024). Can language models trained on written monologue learn to predict spoken dialogue? Cognitive Science, 48(11), e70013. https://doi.org/10.1111/cogs.70013

38 van Manen, M. (2023). What does ChatGPT mean for qualitative health research? Qualitative Health Research, 33(13), 1135-1139. https://doi.org/10.1177/10497323231210816

39 Wan, W., & Low, G. (Eds.). (2015). Elicited metaphor analysis in educational discourse. John Benjamins Publishing Company. https://doi.org/10.1075/milcc.3

Santrauka

Justina Urbonaitė, Yulia Shcherbino.
Žmonių ir didžiųjų kalbos modelių generuojamos smegenų metaforos: metaforų išgavimo tyrimas

Metaforų kūrimas yra veiksmingas analitinis pagrindas, leidžiantis tirti, kaip žmonės ir didieji kalbos modeliai (DKM) konstruoja prasmę, peržengdami pažodinio aprašymo ribas. Šiame tyrime analizuojamos smegenų metaforos, kurias sugeneravo 61 žmogus ir 8 viešai prieinami naudojant vienodas metaforų išgavimo užklausas. Gauti metaforiniai smegenų apibūdinimai buvo suskirstyti į 13 semantinių šaltinio srities kategorijų. Tyrimo rezultatai atskleidė, kad nors abi grupės dažnai rėmėsi technologinėmis metaforomis (ypač konceptualiąja metafora SMEGENYS YRA KOMPIUTERIS), žmonių atsakymai pasižymėjo gerokai didesne semantine įvairove. Žmonės kur kas dažniau naudojo kasdienius, įkūnyta patirtimi grįstus palyginimus, pavyzdžiui, lygino smegenis su kūno dalimis (raumenimis), buities daiktais (kempine) ar gamtos reiškiniais. Tuo tarpu DKM sugeneruotos metaforos buvo abstraktesnės, homogeniškesnės ir dažniausiai orientuotos į sisteminius ar infrastruktūrinius modelius, tokius kaip smegenys kaip tinklas, miestas ar valdymo centras. Be to, DKM pateikti atsakymai buvo vidutiniškai penkis kartus ilgesni ir išsiskyrė itin detaliu metaforų paaiškinimu bei aukštu lingvistiniu sklandumu. Apibendrinant, tyrimas rodo, kad nors DKM geba puikiai atkurti ir struktūruotai paaiškinti konvencionalias metaforas, jų metaforinė kompetencija yra ribota – jiems trūksta šaltinio sričių įvairovės, naujumo ir fizine patirtimi grįsto pagrindimo, kuris yra esminis natūralaus žmogaus metaforinio kūrybiškumo bruožas.

About the Authors

JUSTINA URBONAITĖ

Associate Professor, Vilnius University, Lithuania

Research interests
Cognitive linguistics, contemporary metaphor research, multimodality, (critical) discourse analysis, legal discourse, translation

Address
5 Universiteto str., LT-01131 Vilnius, Lithuania

E-mail
justina.urbonaite@flf.vu.lt

ORCID iD
0009-0003-4735-6792

YULIA SHCHERBINO

BA student, Vilnius University, Lithuania

Research interests
Cognitive science, cognitive linguistics, psycholinguistics, metaphor research

Address
5 Universiteto str., LT-01131 Vilnius, Lithuania

E-mail
yulia.shcherbino@flf.stud.vu.lt

ORCID iD
0009-0006-2760-8894