

Z. KARVALICS LÁSZLÓ

ARTIFICIAL INTELLIGENCE AND LANGUAGE – A DIFFERENT APPROACH TO THE PROBLEM

Abstract

The article explores the intersection of artificial intelligence (AI) and language, critically examining the development and impact of large language models (LLMs) like GPT. It argues that while these models demonstrate significant advancements in generating linguistically coherent text, their achievements are primarily quantitative rather than qualitative. The text highlights fundamental issues such as copyright concerns and the limitations of synthetic data. It also critiques the anthropomorphizing metaphors used to describe AI capabilities, emphasizing the need for precise language to avoid misconceptions about AI's abilities. The article advocates for focusing on AI's practical applications and ethical considerations rather than overestimating its intelligence.

Keywords: artificial intelligence, Language Models, GPT, Metaphor in AI, Synthetic Data, Ethical Considerations in AI

Over the past two years, the success of increasingly large language models (*LLMs*) in artificial intelligence (AI) has brought a growing number of GPT tools, competing with each other and capable of doing more, into the spotlight.¹ GPTs (*Generative Pre-trained Transformers*) are, in simple terms, programs based on neural networks that construct linguistically correct text responses to natural language textual instructions (questions) using their massive text databases on which they have been trained. Moreover, since the elaboration of these answers, thanks to clever underlying algorithms, very often appears to be human-level, the great leap forward in language technology capability has immediately set off a new wave of discourse, even more futile and meaningless than before, about the now human-embattled stage of the growth of intelligence in machine systems.

¹ For the latest chapter in the history of Natural Language Processing (NLP), see ChatGPT Héja (2024).

The double-faced nature of the significant language models

The analysis that gives rise to such high hopes does not consider that the text generation resulting from the increased predictive power of language models (they are better at continuing a sentence they have started) is based on a quantitative, not qualitative, leap in technology. The qualitative gain is in the outputs and the performance. Beyond the “upgrading” of more rudimentary solutions, it is mainly seen where the use of chatbots has become meaningful and possible in areas previously not or only partially covered by automated solutions.

This enthusiasm does not face up to two fundamental strategic problems with the textual raw material reserve: the gradually judicialising copyright concerns about the texts used and the growing threat of a limit due to the running out of textual ‘fodder’. The ‘synthetic data’ escape route (artificially generated texts to teach the model) is dead because the best current solutions are partial and incomplete. Indeed, the existing uncertainties and errors would be multiplied, so the answers must be accepted with many reservations. Behind the misrepresentations, which are often referred to as ‘hallucinations’, lies a lack of appropriate textual content, but worse still, the constructions offered are based on unchecked and ‘non-quality-assured’ content that is inherently flawed, corrupt, biased, erroneous, one-sided, created by irresponsible publishers.

Paradoxically, the very success of this language technology highlights its substantial limitations. It is enough to think that even ten years ago, there were seamlessly functioning systems that could edit fresh sports and stock market data and weather forecasts into complete and flawless news material and thus take over the text production tasks from humans as “robot journalists”. However, this is due to the syntactic preformatting and closed lexicon of the news types in question, with their unambiguous meanings. It is immediately apparent that the system is inefficient and unreliable where these boundary conditions do not exist. However, any other text type or situation where these constraints are met will be capable of being created by a machine element replacing a human one. This is no coincidence: the generation of these text types is mechanical and low-value-added intellectual work, well worth freeing up a valuable lifetime for other purposes.

This is true even for the top performances of GPT systems: translating a technical text from one language to another (more precisely, from many languages to many languages) or modifying the text modality in the same language (producing a summary, having an “academic” style, etc.). However,

what is beyond this, we repeat, seems unattainable with this solution: GPT does *not* manifest *linguistic competence but mathematical-statistical competence*. For the systems that perform operations according to their algorithms, there is no such thing as what gives meaning to the text and the smaller linguistic units that make it up: the world of meanings. The output produced is like a linguistic performance. Still, the system's algorithmic correspondence of the signals produced can only be measured and managed, and the content *carried is entirely insensitive*.

György Csepeli puts it in a much more poetic and tangible form:

“In human consciousness, the infinite outside and the infinite inside meet, made possible by the linguistic competence innate in man. However, competence is only a possibility, actualised by linguistic performance. Generative linguistic models of artificial intelligence do not rely on the linguistic performance of poets and thinkers, the keepers of the shelter of the human found at home in language. The generative language models are based on sentences written or spoken by anybody [...] The generative language models, which select from the enormous and ever-expanding mass of words and sentences already spoken, do not think; they only imitate thinking [...] the language that brings everything close to everyone loses its way to the truth when it is subjected to the dictatorship of the public as the language of anybody. The submerged common language of the machine and man threatens the very essence of man. By imitating human thinking, artificial intelligence opens the door to the technical interpretation of thinking, losing sight of existence as the primary goal of thinking” (Csepeli 2024: 67–8).²

In such a situation, to call two AI systems that send signals to each other *interlocutors* and to consider them revolutionary is a scientific provocation. However, the researchers at the University of Geneva were not afraid to say so, as their experiment showed that their trained language model S-Bert, consisting of 300 million neurons, was able to pass its own learned and verified program to another, *simpler* network so that it could reproduce it (Riverland and Pouget 2024).³ What happened was that instead of calling

² Csepeli expresses this in relation to a thought of Heidegger: ‘thinking that penetrates into being is fundamentally possible in the medium of language’. The infinity of being is followed by that of language: according to Chomsky, language is free to be an infinite set of sentences formed by the rules of grammar, and thus capable of revealing the multidimensional, complex truth of being. How could artificial intelligence, existing in a closed nanovolume of coding with 0s and 1s rather than being, do the same?

³ <https://www.sciencedaily.com/releases/2024/03/240318142438.htm> Pouget seems to forget how one of the ancestors of all chatbots, Eliza, developed between 1964 and 1967, was once “merged” with the parodistic Eliza, a program that offered the illusion of speech

their solution a “two-element prompt cascade”,⁴ they tried to sensationalise it by winking at the scientific attention economy, as if an artificial intelligence had successfully outsmarted an intelligence gargyle.

It is scientifically more correct, but equally meaningless, to call the “linguistic capabilities” of neural networks “*protocol communication*” (Sato et al. 2008) or “*quasi-language*” (Cangelosi and Parisi 1998). How ‘populations’ of neural networks generate a process similar to the evolution of a simple ‘language’ with informative functions in a specific environment undoubtedly has more analogies with natural language change than in a more straightforward machine language processing situation. However, the situation lacks everything that makes natural language a natural language: what *appears to* the observer to be linguistic is, to the machine side, merely the generation of algorithm sequences. Szummer (2014: 41) underestimates the conceptual superficiality of the engineering world: it is not only “*cognitive scientists who tend to mistake analogies for conceptual identity*”.

Linguistic “frontlines”

Meanwhile, the exaggerations, grand pronouncements and the para discourses that build on them (led by machine intelligence that not only reaches but surpasses human intelligence) distract attention from what is happening at other, equally important, intersections of AI and language.

One such area is the “upskilling” of previously developed applications with artificial intelligence solutions to support everyday speech and writing production.

Not only has OCR moved towards the increasingly high-resolution recognition, transcription, digitisation and text-to-text (e.g. with *pen-to-*

understanding, and Doctor, a program that simulated a therapist extracting questions from previous answers. Ostensibly there was a conversation, ‘two agents interacting in a natural language’, but in reality there was a meaningless and self-reversing code multiplication, which as natural language exposed the profound meaninglessness of the whole. In this experiment, of course, there was more than just two new generations of GPT programs being “put together”, as the Betone Studio’s podcast series “Inhuman” does (<https://betone.hu/shows/embertelen/>): in each broadcast, Betti and Peter, “personalised” in different directions, debate current issues. S-Bert has been programmed to further formalize a new “skill” he has acquired during his programming career.

⁴ A prompt is a natural language instruction that is converted into code, processed and executed by the machine system, and the result is presented again in natural language. What happened in Geneva can be seen as a prompt innovation, where the initial instruction included the addition of the ability to form a secondary prompt from the result of the processing.

print tools that reliably convert handwriting into cursive text), but it has also opened up new ways of mass transcription of private diaries, correspondence, old and damaged manuscripts and documents that are difficult to read but of great cultural and historical importance. An imposing development is *In Codice Ratio* (Nieddu 2021) for the Vatican Library, whose creators used a student-based citizen science platform to train the system alongside artificial intelligence.

Solutions that turn live speech into text are also becoming increasingly sophisticated. Some of these also aim to bridge the gap that separates more minor languages, which are under-represented in AI use, from world languages. (For example, a recent development, *Vulavula*,⁵ recognises and describes the spoken names of people and places in four languages spoken in South Africa. Much progress has also been made in the opposite direction, in machine speech synthesis – the 'pronunciation' of written text is so much under siege from natural-looking spoken language that there are reports of new 'robotic speech' systems being put into operation almost weekly. A quite extraordinary best practice is what the community of Maori speakers, led by a non-profit radio station, Te Hiku Media in New Zealand, has done to ensure that they, rather than a global company, are the developers and owners of the speech recognition software that had to be created to digitise their audio archive and the collection of idioms that aids correct identification and transcription.⁶ The project has also renewed the Maori language, but in a way that has avoided the pitfalls of 'digital colonialism' (it is no coincidence that the Mohawk Indians in Canada and the Native Hawaiians have undertaken similar projects).

AI can also be used to research/model several important linguistic issues: the origin of language, emergent aspects of language development and the evolution of language itself (Vogt et al. 2005). AI has also been used to combat language extinctions that have led to the loss of linguistic diversity, most notably in the programme for preserving and conserving African languages.

There seems to be a very long list of issues where language and AI meet, yet we hardly ever address them now in the shadow of ChatGPT. This paper aims to lead an expedition around perhaps the most rarely touched intersection: How appropriate is the use of concepts and language to describe the world of AI and to track its changes?

⁵ <https://beta-vulavula.lalapapa.ai/>

⁶ For details see: <https://blogs.nvidia.com/blog/te-hiku-media-maori-speech-ai/>

Generative schema, cognitive path dependence, metaphor trap

In 1875, John Ruskin satirised the British Metaphysical Society, in which he pens a satirical paper on the widely used buzzword of the day, automation, automatism, the automatic adjective. In the same year, in the columns of *Deucalion*, he marvels, swimming against the tide of the zeitgeist, that patently absurd statements were more easily labelled 'nonsense' in earlier times than in this decade when many people speak with complete naturalism of man being 'an automaton' (Anger 2024: 1–2).

One hundred and fifty years on, and we are pretty much back to square one. "*Actual changes in the working environment and the means of leisure can also give the impression that humans are, after all, just 'biocomputers' that get along well and badly,*" writes Szummer (2014: 41). And of course, the reverse analogy is equally strong: the image and fear of the robotic being competing with the human, and then the position of seeing into the computer, and later into artificial intelligence, a 'thinking machine' with powers beyond the capabilities of the human mind, seems to fundamentally determine the way most people think and talk about this common intersection of technology and human reason (running into Daniel Dennett, who, according to Szummer, doubts the existence of a difference between machine and human intentionality).

Nánay (2000: 22) sees how this situation could have developed.

"The analogy between mind and computer program has long determined, and still partly determines, thinking about the mind. Interestingly, the explanation of the mind and the plan for creating the thinking machine interact. Machines were designed to resemble the human mind, and the human mind was discussed in computer terminology. This two-way effect created a very intricate relationship between the two fields of research, and as computers became more and more complex, this relationship was constantly reassessed."

It does not follow, however, that this conceptual 'crossing over' between machine learning and mental activity is correct and appropriate. In other words, this widespread conceptual and associative teleology is the epistemologically correct choice.

At the same time, it is understandable and accessible to reconstruct how the ideas that drifted towards the world of the automaton and the thinking machine were organised around notions of the organisation of knowledge. In the words of the historian Kurt Danziger, a generative schema, a metaphorical framework, was at work here, too, which, through analogy

through networked categories, made it possible to express perceived changes in reality in language in public dialogue. Whoever becomes embedded in this new web of meanings can draw on it to grasp newer and newer aspects (Anger 2024: 3–7). Moreover, this is one of the basic mechanisms of mind functioning.⁷ To use Lakoff and Johnson's (1981) often-cited basic thesis, metaphor extends our existing knowledge of the world.⁸

So where is the problem? What is wrong with the metaphor cloud built around artificial intelligence?

In short, it has not only *become stultifying*. However, it has also explicitly made objective dialogue difficult, deforming public thinking, influencing strategic decisions about AI, and infiltrating the scientific discourse by deepening conceptual confusion. When it came into being, it played a massive role in ensuring that the most important innovations of the culture-changing technological explosion found their way into understanding and that its developments were easy to imagine and follow, helping to break down the barriers to diffusion. However, what became familiar and entrenched in the early stages of technological take-up has fundamentally influenced how the conceptual web has evolved as AI systems and solutions have become more complex and sophisticated – what I like to call *cognitive path dependency*. Since the machine side was from the very first linked with anthropomorphic associations, the words used to describe each new 'capability' offered a promising solution to the most similar human actions – no longer only in the press, but also in the small worlds of research and development, as terms. Thus, 'machine *learning*' and all its new generations (deep learning, reinforcement learning, federated learning) are based on the well-known notion of the learning process. This is based on the undeniable similarity between the more excellent repertoire of actions created by the knowledge acquired through the behaviour of living beings, known as learning, and the higher operational performance of a machine system assisted by human operators. However, there is a vast difference in several essential attributes. Human learning, whether voluntary or spontaneous, is a form of behaviour that is a kind of intersection of the learning environment, the object of the learning process and the internal state of the learner. It involves a multiplicity of relations of meaning and significance and their relation to

⁷ Draaisma (2005) shows how every technological change since the 17th century has acted like a metaphor machine: offering itself as raw material to refine and deepen our understanding of cognitive processes and memory.

⁸ In a recent paper, Mithen (2024) goes even further: he sees in the metaphorical use of Homo Sapiens the surplus that has been added over Neanderthal man, who did not have this ability.

the external and internal worlds, with choices and actions resulting from them as necessary. Machine learning is not a behaviour but a particular case of program execution, where the output of instructions improves the basic program. You could call it algorithm refinement or model refinement, but it is signal processing in the same way as running a single line of program code. However, each occurrence of the word learning adds to the cognitive path dependency, whereby the anthropomorphic associations with a machine system engaged in mere computation are reinforced, becoming stubbornly inextinguishable. There is no stopping it: just as the machine does not 'think' or 'learn', neither does it 'decide', 'discover' or 'drive' a vehicle.

Artificial intelligence makes no decisions but performs calculations on decision situations modelled by machine code using input data and preformulated processing algorithms. This is true for the most straightforward and most complex systems. The IT mini system that controls the traffic light does not "decide" how many seconds it takes to change from green to amber but executes the program that performs it. Moreover, a stock market program with a vast database does not decide; it only appears *to* decide, but it owes its output to its creators. Decision-making is one of the most complex challenges of the human brain, with many brain areas mobilised simultaneously, depending on the peculiarity, the unusualness, the complexity and the history of the decision situation. The systems called Automated Decision Making (ADM) certainly transfer the decision-making process from human to machine, but not the decision itself: the decision parameters and the decision outputs are "set" by the creators of the systems, and the complex challenges of the decisions are also played out in their brains.

In the same way, AI does not "discover" anything but helps researchers fill knowledge gaps more quickly through its computational power.

At the heart of the 'self-driving vehicle', a complex cybernetic system of sensing, processing and control modules is designed to produce real-time outputs from a real-time data set of vehicle speed, attitude and direction, signal transmission and reception. Instead of moving a physical object in a simulation environment, it moves a physical object in reality. It is *as if* an intelligent entity is driving the vehicle, but all the outputs are determined by the algorithms of the engineering team developing the program. If the control system had an experience, it would not be a driving experience but a computing experience.

In summary, the problem with this metaphorical apparatus of verbs is that it attributes to the notion of action the elements of intention and will, indirectly consciousness and a kind of 'persona', self-consciousness, to

the most advanced programs. This nonsensical and countless times refuted starting point poisons the language built upon it, even through the choice of names. “To speak of ‘artificial intelligence’, in the singular, with a definite article, is valid in a single context: as a research direction, as a scientific-engineering field, as a set of disciplines that can be classified as such.” Thus, of course, it is correct to speak of ‘the’ artificial intelligence research, ‘the’ artificial intelligence systems – but there is no such thing, no such entity, as ‘artificial intelligence’. Many different artificial intelligence worlds of varying complexity and embeddedness share the nature of running software on hardware. These include partial integrations. They can be interconnected with other solutions in the digital ecosystem with other functions: robots, sensors, and virtual reality environments, but the resulting more complex solutions will not be life forms and autonomous entities but will be more similar to them because their functioning will be more reminiscent of sovereign and internally driven action. However, even integrated solutions do not boil down to an all-encompassing giant entity – which can then be alarmed at its dangers and how it is preparing to take over from us.

AI applications are, therefore, neither alive nor behaving: they produce automatic outputs at the end of a processing queue. They do not feel, think, decide, choose, or have no will or intention. They have no consciousness. They are not intelligent. The world of meaning is unknown to them. They process signals according to a preformulated set of instructions. Their outputs are “translated back” into meaning by the human mind that uses them. Contrary to Nietzsche’s oft-quoted words (i.e. that the moving army of metaphors constructs truth for us), it is time to acknowledge that structural metaphors are *‘sometimes, in fortunate cases, constructive’* (Pléh 2003: 405), and to return to the metaphor scepticism of the ancient philosophers. They saw metaphors as interfering with the search for truth and, therefore, potentially damaging to science.

On the other hand, it can be seen that the use of metonymic language, which is also prevalent in the context of artificial intelligence, leads to much less misunderstanding and confusion. For example, ‘AI-ethics’ clearly does not exist, but it is clear to all that it should be understood to mean ‘ethical issues related to the development and application of AI’. We know that AI does not ‘take away anyone’s job’ but can improve the efficiency of production or organisational operations by triggering processes that can replace low-value-added human work with machine operations at a lower cost.

Nevertheless, the part-part relationship, a form of metonymy, is consistently pulled in one direction: the machine Part is always emphasised

in the Whole, composed of engineers, hardware and software components, and the human genius that created it is lost halfway through. This is most visibly reflected in the “machine overcomes man” narrative. A game program called checker was able to beat a human player as early as the late 1950s, and then came chess, go, bridge, and poker – and even though these results were, of course, achieved by a hybrid system of computer scientists and computers that turned the experience of millions of human games into algorithms, somehow they managed to land on the still famous, empty and meaningless ‘machine smarter than man’ turn. Moreover, these exaggerations have given way to severe yet silly scenarios in which the machine, ‘winning’ in a game situation, will, in the knowledge of its superiority, ‘defeat’, ‘replace’ or even ‘exterminate’ the human race itself. Moreover, the shadow of cognitive path dependency looms even behind this extreme, which is ultimately beyond discourse. After all, we are talking about a “form of intelligence”, meaning that even these abstract conclusions are permissible. After all, intelligence implies that everything we already know is not a feature of machine signal processing: consciousness, meaning, complex structures of meaning, will, and intention.

The concept of artificial intelligence as a linguistic “original sin”

In the mid-fifties of the last century, computing applications began to move beyond mere computational power or even elementary data processing, and more and more researchers recognised that it was possible to penetrate further realms of algorithmic brainwork gradually. What can be translated into the language of numbers (bits) – images, fonts, sounds – can be processed in quantity and at a speed that the mind cannot. Moreover, the result of this processing can be translated back into the original form of the content so that machine solutions can support natural information processing well. However, the term “artificial intelligence” (AI), the invention of John McCarthy, was by no means a necessary cover. It lacks the ‘as if’ element, suggesting that the machine is really thinking, not just pretending to think. Interestingly, Nikola Tesla was sensitive to this distinction in his more than half-century-old article for *The Century Magazine* in 1900. He writes of the future “clever” machine (quoted in Colton 2023): Independent of all assistance, left entirely to itself, it will be able to perform several operations in response to external influences on its sensitive components “as if it had intelligence”. Of course, at the end of the 1950s, it was impossible to foresee how the word ‘intelligence’ would be subsumed in the anthropomorphising metaphor

cloud that would not interfere with the developments themselves but would threaten the conceptual clarity of thinking about artificial intelligence.

Nor would the 'complex information processing' proposed by Herbert Simon have been a more appropriate choice. Using it would have created a powerful metonymy trap. After all, computers do not process information; they process signals. Information processing is the exclusive domain of the human brain. However, while we objectify our mental contents to convert them into something the machine can consume, the meaning is lost and only returns when the human brain is confronted again with the result of the signal processing. The contact that is the basis of metonymy undoubtedly exists between the two different qualities,⁹ and it does not interfere with understanding if we use the other name instead of one. However, it cannot be ignored in all contexts where the distinction is significant. It is also the only cure for cognitive path dependence at the moment. In the open communication stream of the Internet, it is difficult to bring about change. However, clarification can be achieved in the academic discourse, and a consensus on the use of terms can be reached gradually. In that case, the arguments and emphases will eventually seep into public discourse and public thinking.

References

- Anger, Suzy 2024. Introduction. The Victorian Automata/Automatism Schema In: Anger, Suzy – Wranks, Thomas (eds.): *Victorian Automata. Mechanism and Agency in the Nineteenth Century*. Cambridge University Press. 1–25. <https://doi.org/10.1017/9781009110129>
- Bencze M. Ildikó 2009. Mi a metonímia? Elméletek és kísérletek a metonímiakutatásban. *Magyar Pszichológiai Szemle*, (64) 4. 677–96. <https://doi.org/10.1556/MpSzle.64.2009.4.4>
- Cangelosi, Angelo – Parisi, Domenico 1998. The Emergence of a “Language” in an Evolving Population of Neural Networks *Connection Science*, 10(2). 83–97. <https://doi.org/10.1080/095400998116512>
- Colton, Emma 2023. Resurfaced Nikola Tesla writings about machines with their 'own mind' eerily predict rise of AI. *Fox News*. May 10. <https://www.foxnews.com/tech/resurfaced-nikola-tesla-writings-about-machines-with-their-own-mind-eerily-predict-rise-of-ai>

⁹ The situation is actually much more complicated. According to Petőfi's (2004:30) typology of text as a complex sign, for example, the marker-signified sign relation actually has six different components that are logically and epistemologically distinct from each other: the physical object (vehiculum), its mental image (vehiculum-imago), the formal organization (formation), the semantic organization (sensus), the factual configuration assumed as a world-part, the relatum, and its mental image, the relatum-imago.

- Csepeli, György 2024. Az új fejedelem. *Educatio* 33 (1). 65–8.
<https://doi.org/10.1556/2063.33.2024.1.7>
- Draaisma, Douwe 2005. *Metaforamasina. Az emlékezet egyik lehetséges története.* Typotex Kiadó. Budapest.
- Héja, Enikő 2024. A ChatGPT története. *Magyar Tudomány* (185)6. 815–28.
<https://doi.org/10.1556/2065.185.2024.6.11>
- Lakoff, George – Johnson, Mark 1981. *Metaphors We Live By.* The University of Chicago Press. Chicago.
- Mithen, Steven 2024. How Neanderthal language differed from modern human – they probably didn’t use metaphors. *The Conversation.* May 20. <https://theconversation.com/how-neanderthal-language-differed-from-modern-human-they-probably-didnt-use-metaphors-229942>
- Nieddu, Elena et al. 2021. In Codice Ratio: A crowd-enabled solution for low resource machine transcription of the Vatican Registers. *Information Processing & Management* (58)5. <https://doi.org/10.1016/j.ipm.2021.102606>
- Petőfi S. János 2004. *A szöveg mint komplex jel. Bevezetés a szemiotikai-textológiai szövegszemléletbe.* Akadémiai Kiadó. Budapest.
- Pléh, Csaba 2003. Douwe Draaisma: Metaforamasina. Az emlékezet egyik lehetséges története. *BUKSZ* (15)4. 404–6.
- Riveland, Reidar – Pouget Alexandre 2024. Natural language instructions induce compositional generalization in networks of neurons. *Nature Neuroscience,* <https://doi.org/10.1038/s41593-024-01607-5>
- Sato, Takashi – Uchibe, Eiji – Doya, Kenji 2008. Learning how, what, and whether to communicate: emergence of protocommunication in reinforcement learning agents. *Artificial Life and Robotics* 12. 70–4.
<https://doi.org/10.1007/s10015-007-0444-x>
- Szumner, Csaba 2014. *Freud, avagy a modernitás mítosza. (Tanulmányok és esszék a pszichoanalitikus antropológiáról és beszédmódról).* KGRE – L’Harmattan. Budapest.
- Vogt, Paul – de Boer, Bart – Belpaeme, Tony 2005. Modelling language origins and evolution *IJCAI 2005 Tutorial* 31 July 2005, Edinburgh, Scotland 1-44. <http://www.ling.ed.ac.uk/paulv/tutorial.html>
- Z. Karvalics, László 2020. Méhek és zebrahalak. Adalékok a “fogalmi bizonytalanság” korához. *Szabad Piac* 2. 47–57.

László Z. Karvalics

Senior Research Fellow

Institute of Advanced Studies, IASK/FTI

E-mail: laszlo.karvalics@iask.hu

<https://orcid.org/0000-0002-3502-434X>