Wittgenstein Ontology: Representing the Philosophical Content of Wittgenstein's Nachlass

Jakub Gomułka

AGH University of Kraków, Faculty of Humanities, ul. Czarnowiejska 36 (C-7), 30-054 Kraków, Poland

Abstract

The paper discusses the efforts to make Ludwig Wittgenstein's *Nachlass* (a collection of all manuscripts, notes, and typescripts left by the philosopher) accessible through the Semantic Web technology. Starting in the 1990s, the Wittgenstein Archives at the University of Bergen (WAB) digitised the collection, leading to its availability on CD-ROM and later online. Today, several WAB's web services that facilitate detailed exploration of Wittgenstein's works are supported by a dedicated OWL knowledge graph called Wittgenstein Ontology. The primary goal of the graph was to represent the whole structure of the *Nachlass* but WAB also aimed at modelling the philosophical content of Wittgenstein's writings. However, the latter task proved to be difficult due to the nature of philosophy itself and has not yet been accomplished. The paper presents certain theoretical solutions, the implementation of which will facilitate the integration of the philosophical content of the *Nachlass* into the existing Wittgenstein Ontology knowledge graph. Furthermore, the paper discusses a proposal for creating a broader system that will allow for the convenient use of this knowledge graph.

Keywords

Ludwig Wittgenstein, knowledge graph, philosophy, OWL

1. Introduction

The present paper discusses the Wittgenstein Ontology project, an effort to create a comprehensive knowledge graph of Ludwig Wittgenstein's *Nachlass* by the Wittgenstein Archives at the University of Bergen (WAB). The project aims to capture and represent not only the structure of the collection of documents left by Wittgenstein but also the complex subject matter of his philosophical conceptions. The latter task is challenging due to the inherent vagueness, inconsistencies, and contradictions within the humanities.

The paper begins by providing an overview of Wittgenstein's works and philosophical contributions. It then examines the development of the Wittgenstein Ontology project. Following this, the paper introduces an approach inspired by medieval theories of supposition to offer a solution for representing philosophical content. The paper concludes with a discussion on the potential role of generative AI in automating the development of the knowledge graph and creating a more accessible and user-friendly interface for the graph.

Proceedings of the Joint Ontology Workshops (JOWO) - Episode X: The Tukker Zomer of Ontology, and satellite events co-located with the 14th International Conference on Formal Ontology in Information Systems (FOIS 2024), July 15-19, 2024, Enschede, The Netherlands.

jgomulka@agh.edu.pl (J. Gomułka)
0000-0002-9100-0334 (J. Gomułka)

2. Wittgenstein's Nachlass

Ludwig Wittgenstein, an Austrian who adopted British citizenship following the *Anschluss*, was an active philosopher from around 1912 until his death in 1951. He is widely recognised as one of the most important thinkers of the twentieth century, and his work continues to provoke numerous controversies and heated interpretative disputes to this day. During his lifetime, he published very little: only one short philosophical book – the *Tractatus Logico-Philosophicus* – two articles, and a German language dictionary for elementary school students. However, he left a vast collection of handwritten notes and typescripts to the executors of his will, based on which several works bearing the philosopher's name were published, including the famous *Philosophical Investigations* [1]. Most of these materials were not publicly available until the 1990s, when the entire collection was first published in print [2], then in electronic form on CD- ROMs[3]. Wittgenstein's *Nachlass*, cataloged by Georg Henrik von Wright[4], one of the executors of the philosopher's will, currently consists of 161 manuscripts, notebooks, and typescripts. Typically, the documents are segmented into paragraphs or remarks, which are occasionally numbered or dated, and vary in length; it is not uncommon for a single page to encompass several of them. The collection comprises nearly 55,000 remarks, encompassing over 185,000 sentences. It is estimated that the entire text could theoretically be printed across 10,000 pages.

3. The creation of Wittgenstein Ontology

In 1990, the University of Bergen established the Wittgenstein Archives (WAB), whose initial aim was to transcribe the entire *Nachlass* into a machine-readable format and make it accessible using dedicated software[5]. In principle, this task was completed by the end of the 90s and in 2000, the aforementioned full electronic edition of Wittgenstein's papers was released[3]. However, WAB continued to work on making the collection available, this time online, within the Semantic Web framework. Between 2006 and 2009, a project was undertaken to create a representation of information concerning the *Nachlass* using RDF and OWL[2]. For this project, a unique ontology with a hierarchy of classes encompassing three branches: Source, Person, and Subject, was created. The first, second, and part of the third of these branches were used to create a knowledge graph including metadata about the structure of the collection: from individual sentences, through remarks, to volumes of manuscripts and typescripts. To this day, the graph, called Wittgenstein Ontology, serves as a backbone for several WAB's web services that facilitate the access to the information about the *Nachlass* (https://wab.uib.no/wittgensteinonline.page), including the direct explorer of the information gathered (http://wab.uib.no/sfb/).

The intentions of researchers associated with WAB were, however, much more far-reaching: the Wittgenstein Ontology was to cover also the content of the philosophical concepts of the author of the *Tractatus Logico- Philosophicus.* The third branch of the class hierarchy, Subject, includes such subclasses as Point and Issue, recently renamed to Claim and Concept; their

[1] The role of the executors extended far beyond mere publishing; their editorial choices are thought to have significantly shaped the image of Wittgenstein as understood by several generations of scholars [1].

[2] It is worth noting that WAB has also played a major role in the development of a semantic manuscript description ontology and a semantic annotation tool within the DM2E project.

instances were to represent what the remarks and documents are about [6, 7, 8, 9]. In their 2021 publication [9], a group of WAB-related researchers discussed reasons why this intention proved difficult to implement: philosophical knowledge is marked by messy conceptual dynamics, sensitivity to changes in context and perspective, systematic underdetermination, inconsistency, and even internal contradictions. To this enumeration, one should also add the observation that philosophical texts typically present, explicitly or implicitly, categorisations of reality, as a result of which they may include different levels of description. A categorisation of this type is offered, for example, by the *Tractatus Logico-Philosophicus*, which once prompted Amélie Zöllner-Weber and Alois Pichler to attempt representing parts of its content as OWL class structures [6]. However, as it turned out, the Tractarian ontology cannot be directly read as an informational ontology in the sense of OWL.

4. Modes of representation

The rather imprecise and provisional recipe for creating philosophical content representation of Wittgenstein's papers included in [9] assumed the creation of instances of the classes Concept and Claim (the latter was still called 'Point' in the paper) based on text mining. Examples of the former included 'logical independence,' 'understanding,' and 'logical analysis'; examples of the latter, in turn, were 'Philosophy requires logical analysis,' and 'Understanding is not a mental process.' Both were to enter into relationships with their kinds, for example, 'understanding' was to be related to its various specifications, namely, such Concept instances, as 'understanding a sentence,' 'understanding a gesture,' or 'understanding music'; in turn, 'Understanding is not a mental process' was to be broken down into the other two Claim instances: 'Understanding is not a process' and 'Understanding is not mental.' Relationships were also to occur between instances of both classes due to the occurrence of given concepts in given claims. The authors emphasised that all these relationships should take into account the ambiguities of philosophical ideas.

Being aware of the discrepancies among the various claims considered in the documents belonging to the *Nachlass*, the researchers envisaged the possibility of creating multiple coherent networks of connections (sub-ontologies). Each network would articulate a distinct viewpoint, represented in the Wittgenstein ontology by instances of the Perspective class. The scholars also proposed adding two other subclasses under Subject: Debate and Argument. Instances of these subclasses would correspond to constructions composed of Claim instances linked by specific properties representing logical relations. As the WAB researchers pointed out, the detection of arguments cannot be fully automated and must rely on human judgment.

Although the article does not mention it directly, it is clear that the future instances of Concept, Claim, Perspective, Argument, and Debate would also have been linked with instances representing individual volumes, pages, or remarks of Wittgenstein's *Nachlass*. In this way, WAB researchers hoped to create a knowledge graph that would have allowed for searching of fragments in which a given piece of content is considered; and, vice versa, checking, what kind of content occurs in specific remarks.

Implementing this proposal would be extremely difficult, and even its authors acknowledged that they were not able to offer a definite answer to the question of how to accommodate both the vagueness and systematic ambiguity of philosophical content with the requirements of formal ontology. Attributing concepts and claims to remarks would be an easy part. However, connecting them through relations to either reveal their internal structure or to construct more complex units would necessitate assuming some set of criteria; after all, the matter cannot rely on the mere intuition of a developer responsible for building the graph. Although, the proposal does suggest that the articulation of claims and concepts would be a responsibility of the developers, moreover, it will be done in English, thereby rendering their association with the original Wittgenstein's handwritings and typescripts – written mostly in German – somewhat conventional in nature.

One can propose an alternative approach, in which the literal content of documents serves as a guide for developing the graph. In this approach, methods of natural language processing (NLP) can be utilised to: 1) decompose sentences into components that will constitute instances of the Claim class; 2) perform grammatical analysis of these components to extract individual words – they will constitute instances of the Concept class – together with their syntactical roles; 3) link Claim instances with Concept instances through relations named after these roles. Claims, in turn would be attached to the instances that represent sentences.

Due to the logical inconsistency of the content of the *Nachlass*, that is, the occurrence of pairs of claims that directly contradict each other [3], representations of the grammatical structures of statements would have to be incorporated not directly, but within the framework provided by the RDF syntax through the mechanism of reification. The use of this mechanism as a response to the problem of inconsistency in philosophical content was proposed earlier [10, 11].

It must be acknowledged that the grammatical solution has a couple of advantages. Firstly, the representation is as close to the text as possible, which essentially eliminates the possibility of dispute over interpretation, since there is almost no room for interpretation at all. Secondly, such a graph can be relatively easily created using available NLP tools. The results will have to be corrected by a human in some cases, but the majority of automatically generated relations will be accurate.

However, a clear drawback of this kind of representation is undoubtedly the content of the Concept class it generates. Its instances would include various forms of articles, as well as pronouns, auxiliary verbs 'sein' and 'haben'; in short, all those essentially insignificant elements of a sentence. Moreover, the grammatical structures of natural language are often ambiguous, both at the level of syntax and semantics. As a result, the application of such an approach would lead to ambiguity within the knowledge graph itself. Consequently, conducting automatic reasoning would be quite difficult.

5. Representing Concepts as occuring in Claims

Theories of supposition developed by medieval philosophers tried to describe how the meanings of words are influenced by their roles within sentences. For instance, in the following three sentences: 'A man runs,' 'Man is a species,' and 'Man is a monosyllabe,' the noun 'man'

[3] It is quite well known that after 1929, Wittgenstein often critically addressed certain theses of his *Tractatus Logico-Philosophicus*. Moreover, many of his remarks are clearly provisional and undermined by later considerations.

has different meanings because it has been used in three different suppositions, respectively: personal, simple, and material [12].

The medieval theories of supposition propose not only interesting but also potentially useful form of natural language analysis. From the perspective of contemporary logic, this pre-quantificational approach can be considered semi-formal, striving to make explicit structures of the natural language text without proposing its translation into expressions of a formal calculus. In creating a representation of philosophical content, on the one hand, we would like to achieve a result that would be as uncontroversial as possible, thus very close to the original text, and on the other hand, we count on the possibility of conducting automatic inferences, thus we aim to establish some form of abstract order. The semi- formal approach offers a reasonable compromise that reconciles these opposing aspirations.

Therefore, the present paper proposes a method of representing the philosophical content of Wittgenstein's works, based on the actual text, from which logical-grammatical structures are extracted to facilitate automatic inferences. In creating this representation, we connect concepts that appear in different roles, examples of which include modes of supposition. For instance, compare the following four use cases of the word 'Gegenstand' in the *Tractatus*:

(a) 'Der Sachverhalt ist eine Verbindung von Gegenständen' (2.01);
(b) 'Der Gegenstand ist einfach' (2.02);
(c) 'Hier entspricht dem schwankenden Gebrauch der Worte „Eigenschaft“ und „Relation“ der schwankende Gebrauch des Wortes „Gegenstand“' (4.123);
(d) 'So ist der variable Name „x“ das eigentliche Zeichen des Scheinbegriffes Gegenstand' (4.1272).

In each of those cases the word 'Gegenstand' occurs in a different role: in (a) we deal with an undefined plurality of 'Gegenstände,' in (b) the reference is to a 'Gegenstand' as a kind, in (c) it concerns a word 'Gegenstand,' and in (d) a concept 'Gegenstand.' In terms of medieval logicians, the first three cases are, respectively, personal, simple, and material supposition.

Within the knowledge graph, the basic form for representing any concept will be the ordinary entity of the Concept class. This initial form will correspond to the simplest roles: in the case of nouns, it will function in the mode of simple supposition, and for verbs, the indefinite tense (represented by the present tense, singular, third person). Further, to represent other roles, we expand the ontology with subclasses of the Concept class like ConceptVariable, Set, and Quantification. Each entity belonging to one of the subclasses simultaneously belongs to the Concept class and must be associated with one of the ordinary entities. The name 'Quantification' may hint at the creation of first-order calculus constructs. However, in reality, this subclass is intended to cover more cases than just the equivalents of the universal and existential quantifiers.

Let us examine the aforementioned example from the *Tractatus*: 'Der Sachverhalt ist eine Verbindung von Gegenständen.' The name 'Gegenstand,' as we said, is used to refer to an undefined plurality of its designates. In order to express this one needs to create an instance of Quantification class; that is, a quantified variable of a basic form 'Gegenstand.' Next this Quantification instance is attributed with a special data property 'plural' which marks its plurality. Since Quantification is a subclass of Concept, that instance is also a kind of Concept that can occur as a complement in a complex concept 'Verbindung von Gegenständen.'

Take the next example sentence from the *Tractatus*: 'Die Welt ist alles, was der Fall ist.' Here, we have a concept 'Fall sein.' This concept is used in a quite specific way: the entire phrase 'alles, was der Fall ist' denotes one complex entity composed of individual instances, each of which being the case. This represents a rather complex syntactic structure that requires assuming a basic form that is not present in the sentence: 'Etwas.' The word 'Etwas' appears multiple times in the *Tractatus*, so this is not completely artificial. Since it has no technical employment in the *Tractatus*, it can be used to denote anything, any individual. So, on the concept 'Etwas,' we build an instance of Quantification with the data property 'general,' indicating that it refers to all designates simultaneously. The final layer of this construction is an instance of the class Set (also a subclass of Concept), to which, through the object property 'quantification,' this instance of Quantification is assigned, and also, through the object property 'satisfies,' the concept 'Fall sein.' In other words, the set is defined as containing all designates of the concept 'Etwas' – essentially all individuals – which satisfies the predicate 'Fall sein,' making it a construction equivalent to the expression 'alles, was der Fall ist.'

In creating the knowledge graph, we adopt the convention, according to which the base instances of the Concept class are given names in German. Meanwhile, the properties that connect these instances and are responsible for creating semi-formal structures are named in English. This can be demonstrated using the example of the first proposition of the *Tractatus* mentioned earlier: 'Die Welt ist alles, was der Fall ist.' The proposition is mapped into a structure that includes three basic instances of the Concept class: 'Welt,' 'Etwas,' and 'Fall sein.' Furthermore, it includes one Quantification entity and one Set entity: their get their technical names in English, that is, respectively, 'Quantification W-TLP: 1 [1][1]' and 'Set W-TLP: 1 [1][1].' In the previous paragraph, it is explained how the two are connected to the basic concepts 'Etwas' and 'Fall sein.' Finally, the proposition is rendered to a triple whose subject, property, and object are, respectively, the base concept 'Welt,' the structural property 'equivalent,' and the set 'Set W- TLP: 1 [1][1].'

The *Tractatus* itself is a work that is relatively coherent logically (though it is not without certain inconsistencies), so in principle, one could be tempted to enter representations of its propositions into the knowledge graph directly. However, if our goal is to create a representation of the entire *Nachlass*, we cannot do this and must use the reification mechanism mentioned in the previous section. This move introduces additional complexity to the graph and significantly increases its volume but it seems inevitable.

As work progresses on generating the knowledge graph, the set of structural devices will need to be gradually supplemented to enable the mapping of all content written in natural language. This particularly applies to structures responsible for indicating the proper role of verbs (grammatical person, number, tenses). Some verbs, especially 'sein' (to be) and 'haben' (to have) in their zero-tense usages, as is the case in the aforementioned example, should be mapped as structural properties.

6. User interface: a broader system

The complexity and user-unfriendly nature of the representation structure sketched above is immediately evident. Creating SPARQL queries that provide meaningful insights will be feasible,

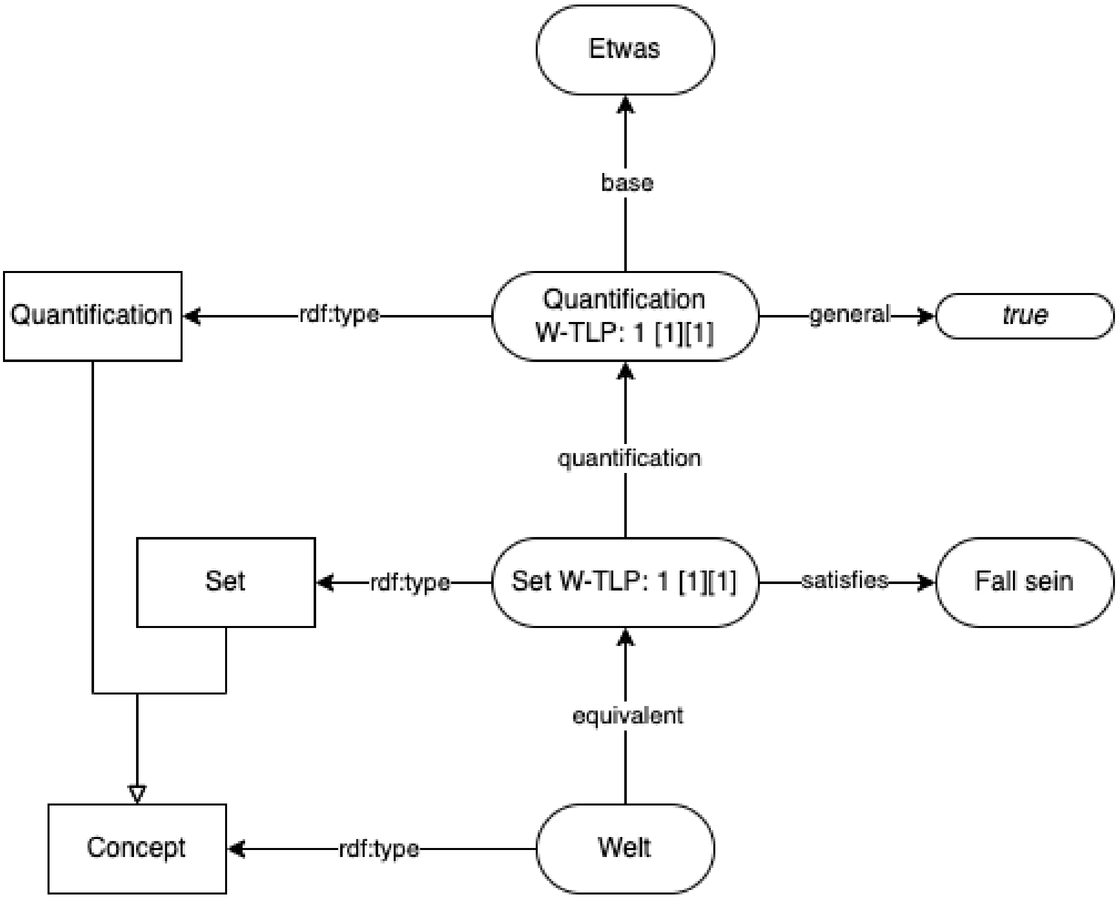

Figure 1: a graph representation of the first thesis of the *Tractatus Logico-Philosophicus*.

but such queries are likely to be extensive and difficult to understand. Consequently, there is a need to create a system that will to act as a sophisticated interface for such a knowledge graph. While it's beyond the scope of this paper to delve into the details of this system, it can be said that it aims to incorporate components that operate as classical algorithmic processors with clearly defined rules for generating and interpreting SPARQL queries, along with a module that employs large language models (LLMs) technology to facilitate natural language interaction with users.

The target system would allow for asking questions in natural language, to which it would provide answers consistent with the knowledge stored in the graph. It's important to emphasise that the LLM itself would serve as a translator of colloquial speech into a format interpretable by the SPARQL query generator and would in no way play the role of the source of information conveyed to the user – this role would be entirely attributed to the knowledge graph. In this way, it would be possible to avoid the phenomenon of AI hallucinations, i.e., the provision of false information by the LLM.[4]

[4]Hallucination phenomena affect even customised GPTs created by subscribers of the ChatGPT service offered by

LLM technology can also be applied in automating the development of the *Nachlass* knowledge graph. A manual transcription of nearly 200,000 sentences into RDF triples can be streamlined with dedicated tools and parsers that would substantially increase the efficiency of the process; and yet, it would be a task requiring enormous amounts of work, perhaps, comparable to the efforts spent on the transcription of Wittgenstein's manuscripts carried out by WAB in the 1990s. Further research should determine whether it will be possible to use a standard available language model, such as ChatGPT 4.0 or most powerful Llama clones, by constructing an appropriate prompt, or whether a separate model will need to be trained.

Developing an effective workflow for Wittgenstein's *Nachlass*, both in terms of creating a knowledge graph and a dedicated user interface, could serve as a starting point for creating similar knowledge bases on other topics. The same or slightly modified semi-formal model of content expressed in natural language could be applied to other philosophical concepts, and likely also in other branches of the humanities. It may not always make sense to use a knowledge graph: this depends on the extent to which automatic reasoning mechanisms can be implemented and the internal coherence of a given content. However, the proposed integration of knowledge graph technology with LLM technology could itself represent a step towards developing an artificial intelligence system that yields reliable information as well as addressing the problem of AI hallucinations.

Acknowledgments

I am very grateful to the University of Bergen and the Wittgenstein Archives for inviting me to a research visit in February 2024 and for creating excellent conditions for my work. In particular, I would like to thank the director of WAB, Alois Pichler, for his invaluable help and support. I also owe words of thanks to Simo Säätelä, Jan Ole Bangen, Filippo Mosca, Andreas Lothe Opdahl, Krzysztof Kutt, Claus Huitfeldt, Heinz Wilhelm Krüger, Rune Jensen Falch, Tim Grasshöfer, as well as to all other participants of the seminars organised by WAB during my stay for their illuminating remarks and questions.

References

[1] T. Wallgren (Ed.), The Creation of Wittgenstein : Understanding the Roles of Rush Rhees, Elizabeth Anscombe and Georg Henrik von Wright, Bloomsbury, London, 2023.

[2] L. Wittgenstein, Wiener Ausgabe, Springer, Vienna, 1994.

[3] L. Wittgenstein, Wittgenstein's Nachlass. Text and Facsimile Version. The Bergen Electronic Edition, Oxford University Press, Oxford, 2000.

[4] G.-H. von Wright, The wittgenstein papers, The Philosophical Review 78 (1969) 483–503.

[5] C. Huitfeldt, Toward a maschine-readable version of wittgenstein's nachlaß: Some editorial problems, in: H. G. Senger (Ed.), Philosophische Editionen: Beiträge zur VI. Internationalen Fachtagung der Arbeitsgemeinschaft philosophischer Editionen (11.–13. Juni 1992 Berlin), Max Niemeyer Verlag, Tübingen, 1994, pp. 37–43.

OpenAI, despite the fact that model customisation involves providing it with a file containing explicit information [13].

[6] A. Zöllner-Weber, A. Pichler, Utilizing owl for wittgenstein's tractatus, in: H. Hrachovec, A. Pichler, J. Wang (Eds.), Papers of the 30th International Ludwig Wittgenstein Symposium (5–11 August 2007, Kirchberg am Wechsel), ALWS, Kirchberg am Wechsel, 2007, pp. 249–251.
[7] A. Pichler, A. Zöllner-Weber, Sharing and debating wittgenstein by using an ontology, Literary and Linguistic Computing 28 (2013) 700–707. doi:10.1093/llc/fqt049.
[8] J. Mácha, R. J. Falch, A. Pichler, Overlapping and competing ontologies in digital humanities, 2013. doi:10.1145/2517978.2517984.
[9] A. Pichler, J. M. Fielding, N. Gangopadhyay, A. L. Opdahl, Crisscross ontology: Mapping concept dynamics, competing argument and multiperspectival knowledge in philosophy, Quaderni di "Filosofia" 2 (2021) 59–73.
[10] J. Gomułka, Artificial intelligence applied to philosophy: A contribution to the wittgenstein ontology project, Ethos 36 (2023) 207–234. doi:10.12887/36-2023-3-143-12.
[11] J. Gomułka, Towards a computational ontology for the philosophy of wittgenstein: Representing aspects of the tractarian philosophy of mathematics, Analiza i Egzystencja 63 (2023) 27–54. doi:10.18276/aie.2023.63-02.
[12] P. V. Spade, Thoughts, words and things: An introduction to late mediaeval logic and semantic theory, 2002. URL: https://www.pvspade.com/Logic/docs/Thoughts,%20Words%20and%20Things1_2.pdf.
[13] F. Mosca, J. Gomułka, What wittgensteinian gpts can't do, forthcoming.