

University of Tartu
Department of Semiotics

Andrea Barone Renolfi
Modeling Futures: On Language Models and the Social Construction of
Knowledge in Academia
Master's Thesis

Supervisor: Ott Puumeister PhD

Tartu

2025

Annotation

This thesis explores the impact of large language models on the social construction of knowledge in academia. Grounded by an actualized present as perceived by specialized and non-specialized researchers, it utilizes the three moments of the dialectic as presented by Berger & Luckmann to compare five problematic areas for AI implementation in academic processes. A comparison takes place between the possible future where AI permeates academic processes and the same processes as they are currently performed. The thesis employs speculative semiotics to construct a sociotechnical imaginary that frames and informs the analysis. The result is an exploration of the possibilities of AI in academia and the consequences of its implementation on the construction of knowledge in academic processes.

Keywords: artificial intelligence, large language modes, academic processes, social construction of knowledge, speculative semiotics, speculative research.

Table of Contents

Annotation	2
Table of Contents	3
Introduction	4
1 AI's Recent Present and Potential Future	7
1.1 LLMs Current Capabilities and Limitations	8
1.2 LLMs Perceived Potential Applications	12
1.3 LLMs in Academia	14
2 Speculative Semiotics as a Technique	18
2.1 Overview of Speculative Research	18
2.2 Speculative Semiotics Next Steps	21
2.3 Speculating on AI and Academia	24
3 Knowledge as a Social Endeavor	30
3.1 Externalization	31
3.2 Internalization	32
3.3 Objectivation	33
3.4 Language and the Construction of Reality	34
4 Analysis of a Becoming of the Present	37
4.1 On the Possibilities of Fetching Data	38
4.2 On the Processes of Writing and Editing	42
4.3 On Interlingual Knowledge Sharing	47
4.4 On Interdisciplinary Research Efforts	51
4.5 On the Epistemological Concerns of Outputs	53
Discussion	56
Conclusion	58
References	61
Kokkuvõte	63

Introduction

As artificial intelligence's ubiquity increases, questions regarding its capabilities, limitations, and effective applications continue to arise. This would be an obvious statement if it were referring to specialized researchers, as they are bound to question the current state and consequences of their own object of study. But these inquiries are also happening outside the specific research community as curiosity about AI's possible implementations bleeds into other areas of knowledge. The *what is* and *what could it be* of AI is present in everyone's minds as we race to catch up with yet another technological development. But there are also cautionary voices present that remind us that we should always be mindful of unintended consequences and possible ethical concerns. Especially as AI starts to get inserted into academia and research, and sediments its place as part of the system. While some proponents exalt the multitude of benefits it could bring, others are more reticent about embracing this new technology. As such, the discourse surrounding artificial intelligence (AI), natural language processing (NLP), large language models (LLMs), and more is continually evolving and incredibly diverse. This has created an environment in which versions of possible futures compete to dominate the current discourse on AI use. Academia in particular, being the bastion of objective knowledge construction and one of the main legitimating forces of social objectifications, needs to be extra cautious when implementing any technological development that might affect the epistemological standing of its productions. And while there have been numerous studies that deal with the current applications of LLMs in different facets of academic practice, few deal with the consequences of the possibility of AI permeating different aspects of the academic knowledge construction processes and the impact this could have on the standing of said processes and their outputs.

In this thesis, I will attempt to construct one particular sociotechnical imaginary, that of a society in which AI has permeated academic processes at different stages. A sociotechnical imaginary will be understood as stipulated by Jasanoff & Kim as "collectively held, institutionally stabilized, and publicly performed visions of desirable futures, animated by shared understandings of forms of social life and social order attainable through, and

supportive of, advances in science and technology” (2015). However, less emphasis will be placed on the importance of the word ‘desirable’ as opposed to the book “Dreamscapes of Modernity” where the definition comes from. This is due to the necessity of finding a problematic present situation in which to ground the speculative process, as it will be explained later. Nevertheless, in order to achieve this sociotechnical imaginary, I will start by establishing the current discourse surrounding LLMs in academia, both by specialized researchers and by non-specialized researchers. As such, the sources presented in chapter one are not intended to be assessed by their empirical value but as a showcase of the different discourses surrounding LLMs in academia from different stakeholders, including researchers working actively in the development of this technology and those that represent oppositional forces. By providing a variety of points of view, chapter one will then serve as the ground for building this sociotechnical imaginary onto which speculative research will be applied. This latter one will be explored in more detail in chapter two. Speculative research, and in particular speculative semiotics, will provide a framework for the projection of a possibility of becoming of a present represented by the sociotechnical imaginary and grounded by the literature showcased in chapter one. In chapter three, the particular socio-semiotic framework that will be used to analyze the sociotechnical imaginary will be laid out, this will mainly consist of Berger and Luckmann’s manifesto on the sociology of knowledge by the name “The Social Construction of Reality” (1966). The analysis itself of the constructed imaginary future will be performed in chapter four.

The methodological choices made in this thesis, particularly that of the use of speculative semiotics, will allow a type of analysis that would be difficult otherwise. A sociotechnical imaginary is sedimented through social institutions and held through a shared understanding of advances in science and technology (Jasanoff, Kim 2015). As such and in virtue of being “visions of desirable futures” (Ibid), a technique that creates such an imaginary world based on current discourses that are used as a ground for their conception is needed. Speculative research as a technique allows precisely that, the creation of a possible world, grounded on current discourses on a particular topic, which can then be analyzed through a semiotic model giving way to a speculative semiotics approach. However, it is important to keep in mind that this approach, while allowing a conversation to occur in the present through

a common understanding of a possibility of becoming of said present, also presents certain restrictions that need to be taken into account. Mainly that this approach, lacking an actualized reality to stand upon and presenting itself as an exercise that attempts to affect the transition into different futures, cannot speak about current realities as they stand. As such, there is no definitive question, in the form of a falsifiable hypotheses, that this thesis seeks to pursue nor it attempts to provide any concrete answer. Instead it is exploratory in nature, looking to explore the possibilities of: How could large language models impact the social construction of knowledge in academia? It merely intends to contribute to the ever-expanding discourse of possibilities of the future and to serve as a medium for the reader to communicate with their own present through its potentiality of becoming and thus affect the possibilities of our own making.

1 AI's Recent Present and Potential Future

While AI may be the acronym in everyone's mouths, artificial intelligence as such refers more to the capacity of machines or computers to mimic human intelligence. In and of itself, it is not the system as we perceive it, but a general label geared towards describing a common objective in a specific field of computer sciences, that is *artificial intelligence*. For the purposes of this thesis, more specific terms will be needed at certain points, however, 'AI' will continue to be used as a general term when technical terms are not specifically required. This choice has the intention to avoid burdening the reader with an overabundance of specialized acronyms when 'AI' has already been incorporated into the common lexicon of non-specialized researchers and is amply used in the literature interchangeably when speaking about large language models. Given that the objective of this thesis is to engage with the process of knowledge construction in academia and also because of the prevalence of AIs such as GPT, Gemini, Deepseek, and others, the scope of possible AIs will be limited to language models under the label 'natural language processing' or NLP. These latter ones have evolved substantially over the years with their more recent developments being large language models or LLMs as they have quickly gained the attention of developers and users alike. While this thesis is not particularly concerned with technical nuances, it is important to determine which area within AI we will be dealing with. Given the amount of attention that LLMs have received in publications and funding alike, it is clear that it is being perceived as having enormous future potential (Ganguli et al. 2022). This thesis will concentrate on LLMs in particular, not only because of their wide range of applications but also because this is the type of AI that academia is already incorporating into its daily processes.

In this chapter, I will provide a brief overview of LLMs as understood by specialized researchers including their current capabilities and limitations as well as their potential present and future applications. In section 1.3 I will broaden the scope and include various publications from different fields and their expectations and reservations regarding AI in academia. Because this is not my area of expertise I have limited the scope of this overview to peer-reviewed papers or highly legitimate sources. Given the lengthy procedures needed for

publication, most of them are at least a year old. Due to the rapid development of the field, it could be considered outdated information. However, this will not be an issue because of the nature of the analysis of this thesis which will be explained in chapter two. In this chapter I mainly aim to establish a baseline of the state of AI at a given point, that is the recent present, and its expected potentiality which, in the current environment of high expectations is rich and plentiful, albeit sometimes somewhat cautious.

1.1 LLMs Current Capabilities and Limitations

Large language models are one of the latest developments within the field of natural language processing which in itself is a subfield of artificial intelligence. Its predecessor, that is language models or LMs, “...refer to systems which are trained on string prediction tasks: that is, predicting the likelihood of a token (character, word or string) given either its preceding context or (in bidirectional and masked LMs) its surrounding context” (Bender et al. 2021: 611). The capabilities of LLMs are, as the acronym NLP would suggest, those associated with natural language queries and differ from LMs by their size which “have been scaled primarily along three factors: amount of computation, number of model parameters, and training dataset size” (Wei et al. 2022: 2). LLMs capabilities have been extended due to scaling and consist of but are not limited to: text generation, question answering, language translation, and text classification (AlZu’bi et al. 2023: 30). Some of the options in the market are purported to be very versatile in their generalization ability and can respond to a wide variety of inquiries and requests with a high level of precision (Ibid). In addition, LLMs differentiate themselves from older models by their increased scale, context awareness, and text generation ability allowing them to respond in a more natural manner mimicking the tone and style of human speech, be less repetitive in their interactions, and have the capacity to respond to more ambiguous queries (Koubaa et al. 2023).

Among the most popular ones exists GPT, that is, Generative Pre-trained Transformer, which surpassed 400 million users in February 2025 (Kant 2025). *Generative* refers to the capacity of the AI to create new data and not merely repurpose data it already possesses, or at

least make it seem like new data based on the level of complexity of the algorithm in question (Zewe 2023). The word *pre-trained* comes from the previous generation of AIs and refers to the data that informs the model, in the case of LLMs and GPT this is the step in which scale turns a pre-trained model into a large language model (AlZu'bi et al. 2023: 30). The word *transformers* refers to a specific technological development that allows the AI to establish dependencies between elements and establish the relative weight of the different tokens (Koubaa et al. 2023), that is, it allows the AI to connect some elements to others and to determine which connections and which elements are more important than others. So in other words, GPT is a large language model at its core, pre-trained with an incredibly large set of data, designed with a high level of complexity to be able to create “new” data, and employing transformer technology which allows it to connect different elements to each other and rank them by level of importance. It ports all the latest developments in AI technology in one single model which is available to the public.

The remaining element to be explored is the second-order significance of the size of large language models. Most of the consequences of increasing the scale of language models can be predicted to some degree. In particular, the improvement of performance that comes from an increase in scale can be predicted with sufficient accuracy (Wei et al. 2022: 1). But some abilities are present in large-scale models that were sometimes completely absent in small-scale models and appear suddenly at certain thresholds, these are called *emergent abilities*. They are named after the concept of *emergence* used in several disciplines that refers to “when quantitative changes in a system result in qualitative changes in behavior” (Ibid: 2). Several emergent abilities have been recorded in modern LLMs such as analogical reasoning, chain-of-thought reasoning, instruction following, explanation leveraging, and many more (Webb et al. 2023; Wei et al. 2022). While the exact mechanism of how these emergent abilities arise is not important for this thesis, their existence and the unpredictable nature of their appearance are two important characteristics that will be brought up in the following chapters. It is also critical to note the existence of *emergent risks*. In the same way that arguably useful abilities, such as reasoning abilities that might surpass human levels have developed in LLMs (Webb et al. 2023), other axes of emergence include truthfulness, bias, and toxicity (Wei et al. 2022: 8). This means that the significance of emergence is not limited

to the abilities of the AI itself to imitate human thought processes, they can also develop and showcase traits that mirror some of the more unsavory sides of human communication and interactions. This is mainly due to the unequal representation of different demographics in the unfathomable large quantity of data that is being used to train the models, the data mainly favoring certain languages and privileged social groups, particularly young people from developed countries (Bender et al. 2021: 613).

A caveat to the idea of emergent abilities arises from the insights of Bender et al. where they repeatedly make the distinction between natural language processing and *natural language understanding*, that is NLU. In the text written in 2021, they explain that languages are systems of signs that pair form and content but that LMs have no access to content nor meaning, and can only process form. They also mention that Karen Spärck Jones wrote in 2004 about the epistemological underpinnings of language and information processing, that is LIP, used in LMs, and said the following:

“...the use of LMs ties us to certain (usually unstated) epistemological and methodological commitments. Either i) we commit ourselves to a noisy-channel interpretation of the task (which rarely makes sense outside of ASR), ii) we abandon any goals of theoretical insight into tasks and treat LMs as “just some convenient technology,” or iii) we implicitly assume a certain statistical relationship — known to be invalid — between inputs, outputs and meanings. Although she primarily had n-gram models in mind, the conclusions remain apt and relevant.” (Bender et al. 2021: 616).

Spärck Jones (2004) on their part seemed to be concerned with the underlying framework of using LMs for LIP at the time. The epistemological and methodological commitments mentioned by Bender et al. seem to be discussed under the umbrella of a criticism of the “deep generation” type of model which according to them “might be seen like ruthless pragmatism” (Ibid: 7). And even though Spärck Jones proceeds to offer a different framework, that of a “deep distributional” type of model, it is difficult to say whether this new vision truly gets rid of the previously stated epistemological commitments of LM use, hence Bender’s et al. reiteration of Spärck Jones’ conclusions remaining relevant.

There are other limitations associated with LMs in general that do not necessarily come from their large counterparts. Some of these issues are due to the nature of what these models are and what they can do, and what humans perceive their capabilities to be. As it has been established LMs process language as tokens and based on previous inputs predict the most likely next outcome. In this way they are able to process language but under no

circumstances are they capable of understanding it, that is, LMs are capable of NLP but not of NLU (Bender et al. 2021: 615). However, humans have the propensity to assign intent to communicative processes so even though LMs are concerned purely with form and not with content, humans that engage with LM's outputs have the propensity to assign motivation to seemingly linguistically fluent language attempts (Ibid: 615-617). This is merely due to human proclivities that can eventually lead to negative real-world applications if the outputs of LMs are taken to be entirely truthful and meaningful, from wrongful arrests due to failure of machine translation to the perpetuation of harmful points of view that can lead to the continued discrimination of different social groups to different degrees (Ibid: 617). In addition, the fluency of a certain output can lead to a higher certainty of accuracy, which is the case with modern machine translations which are perceived as accurate due to their high level of fluency and coherency despite being inaccurate on a regular basis (Ibid: 617-618). All of these instances described so far involve the deployment and use of LMs with no malicious intent, which will be the assumption throughout this thesis.

There have been a multitude of suggestions for how to mitigate these risks both as they relate to the development procedures as well as on a social and governmental level. Among these, we find:

- Calls for curation of data sets and documentation through manual curation to avoid pitfalls such as those found through keyword filtering (Bender et al. 2021: 614).
- Technical mechanisms to avoid data stagnation in large models and bias codification (Bender et al. 2021: 615).
- The creation of documentation that specifies the design objectives for each model, what they are tuned for, and why specific data was selected for training; including the identification of stakeholders to allow users to better understand the frameworks of each model (Bender et al. 2021: 618).
- The implementation of pre-mortem exercises to force development teams to analyze hypothetical failures and reverse engineer solutions to increase the awareness of potential risks (Bender et al. 2021: 618).

- To reduce the asymmetry in compute capacities between the private sector developing the technology and academia, allowing public institutions to analyze potential inherent problems (Ganguli et al. 2022: 1756).
- Implementation of the *red team* method often employed in computer security to discover potential harms prior to model deployment (Ganguli et al. 2022: 1756).
- Implementation of “novel governance structures and government interventions” which could implement legislation that alters the incentives of private sector developers (Ganguli et al. 2022: 1757).
- Improvement of prompting techniques to improve our understanding of how inputs affect outputs (Wei et al. 2022: 11).

1.2 LLMs Perceived Potential Applications

As the use of AI, particularly of LMs and LLMs, becomes widespread suggestions arise for more and more possible applications of these models in everyday life. These range from facilitating human tasks to automatization and productivity improvements to outright creative endeavors, and they span many domains of human activity. While myriad applications for AI have already been realized, in this section I will briefly delve into the different realms of potential AI applications as perceived by researchers. This will help in giving a view of how AI capabilities can be applied and in which direction development might go. It is important to note that none of the applications mentioned in this section are discussed thoroughly nor evaluated, they are merely listed to provide a wide perspective of possibilities.

Businesses and the corporate world are usually the first ones to adopt new technologies if they perceive that the outcome could result in increased productivity and thus profits. AI is not different in this regard. LLMs have been employed as virtual assistants for some time now and in the automation of particular tedious tasks in programming settings (Wei et al. 2022: 32). This is without considering modern industrial applications in customer service and marketing where it can improve efficiency, customer satisfaction, and provide valuable insights (Koubaa et al. 2023). But AI could do so much more, they say. Some of the

purported benefits of letting AI into your life are: having a personal assistant, language and writing aid and translation, providing the means of a creative boost and inspiration, entertainment purposes, and study assistance (Wei et al. 2022:32). But this is not the limit, as AI can also help you make better use of your assets and aid you in your career. LLMs can aid in investing by providing market insight and help you understand market trends by aggregating enormous quantities of data, they can help you with content generation, copywriting, social media management, interview rehearsals, and even research assistance (Ibid). It seems like AI can do it all, at least according to its more ardent proponents. Most of these purported applications are at the individual level and they speak about the perceived possibility of a deep impact of AI in the everyday life of human beings.

Health is another area in which AI use has been researched and considered, particularly as it refers to the healthcare industry. Koubaa et al. (2023) write that “ChatGPT has the potential to revolutionize the healthcare industry by improving patient outcomes, reducing costs, and facilitating more efficient and accurate diagnosis and treatment.” They proceed to give a rundown of different possible medical applications, from everyday uses such as simplifying radiology reports or providing ophthalmology consultations at low costs with sufficient accuracy all the way to the possibility of personalized medicine and the potential of AI-powered data analysis for medical decision-making (Ibid). In addition, Koubaa et al. (2023) also mention previous studies on the possibility of using AI to provide medical diagnosis in certain parts of the world that lack the necessary institutional infrastructure. However, among all the studies mentioned by Koubaa et al. (2023), there is a common missing element, sufficient discussion concerning the ethics of the application of AI and LMs in any healthcare-related area. However, there are others who have a wholly positive outlook for the implementation of AI in healthcare. Khan exalts the potentiality of AI to bring increased objectivity to medical practice and an obviation of biases, they believe that “superior AI generations will diagnose, operate, monitor patients with better accuracy and most importantly, with precision, [...] it will be able to exactly define individualized pharmacogenomics for the best medical outcomes” (2023: 1085). All of this stemming from their belief that “human cannot accept human error” (Ibid: 1086) and these machines need to fill that place under strict regulatory interventions and policies. But again, neither of these

optimistic approaches sufficiently discusses or touches upon the ethical concerns of, as Koubaa et al. put it, “sufficient accuracy” (2023).

Education has seen some of the most striking developments and has had a high degree of discussion regarding its ethical concerns. Koubaa et al. (2023) mention several studies that have suggested AI be used in a plethora of areas. The capabilities of LLMs have been assessed when it comes to the generation of mathematically valid statements, their use as a library reference service, and even their high level of medical knowledge which would render it a useful assistance in medical training that could provide personalized assessments and lessons (Ibid). Some practical examples of AI applications in educational settings can be found in Pinzolit's paper where they mention not only several benefits for students improving engagement and motivation through possible gamification practices, but also provide benefits for teachers by reducing workloads by aiding in student assessment, grading, and even providing feedback (2023: 42-43). In addition, there have been studies that show the possibility of utilizing LLMs to create texts that are indistinguishable from those produced by humans jeopardizing the legitimacy of online exams and there are increasing concerns regarding academic integrity when it comes to the use of AI (Koubaa et al. 2023). Many calls have been made to increase teacher competencies and skills in AI prior to any possible application of these systems in real-world educational settings (Pinzolit 2023: 43). This is the only type of application where mentions of ethical concerns represent a large part of the literature as opposed to others where the plausibility of application and ethics are largely omitted.

1.3 LLMs in Academia

Given the topic of this thesis, particular emphasis needs to be placed on the relationship between LLMs and academia. And there is a clear prevalence of a specific view when it comes to non-specialized researchers. That is, AI is a tool. In this section, I will briefly explore the different views held by academics that do not specifically research NLP. I

will attempt to showcase the expected uses of AI in different fields and in higher education and also any concerns that they might have regarding its implementation.

Pinzolits (2023) makes two assertions that are salient in this discussion. Repeated mentions of AI as a tool throughout their 2023 paper and outright statements of AI and particularly NLP systems having the ability to “understand, interpret, and generate human language” (Ibid: 38). As it has been previously discussed this is patently impossible given the nature of NLP systems that simply deal with predicting token placement. However, Pinzolits assertions provide adequate backing to previous mentions of users mistaking outputs for communication attempts, this type of interaction between users and LLMs will be of great significance throughout this thesis. Other interesting assertions include that in academia AI has led to the development of “tools for literature search, content analysis, scientific writing, and editing” and that these tools promise to “add depth and insight to the content of academic studies” (Ibid: 45) with the cautionary tale that “AI tools influence the stylistic forms of scientific writings” (Ibid: 46). Though granted Pinzolits recognizes the need for academics to be cautious and separate the tool from their own knowledge and creativity, to allow the AI to complement rather than replace their efforts; including the call to implement an ethical framework for AI (Ibid: 42). They also mention the need for “AI governance to move from principals to processes” (Ibid: 42) though they fail to mention what shape these processes would potentially take, a common shortcoming in this type of call to action.

When it comes to practical applications some of the more commonly mentioned are: literature search, analysis of research articles, and academic writing and editing. Pinzolits (2023: 43) emphasized that AI “utilize semantic search capabilities to comprehend the researcher’s intent and provide more relevant results based on the context of the keywords” (2023: 43) a similar sentiment to that expressed by Livberber & Ayvaz (2023). This would entail that some sectors in academia believe that AI has the ability to understand intent based on the words chosen by the user, which again would mean attributing NLU to LMs. In addition, Pinzolits (2023: 44-46) also mentions AI’s ability to deepen analysis by streamlining the literature review process, analyze sentiment through tone, detect bias to protect research integrity, aid with compliance with academic writing standards, provide assistance to non-native speakers in language use, and more. Some of these possibilities have already been

undermined by the previous sections, particularly those related to assistance to non-native speakers and other low-resource language users and those related to biases.

Some authors have advocated for the use of AI to help with repetitive or monotonous tasks. Koubaa et al. (2023) mention the opportunity that the development of AI in academia could bring when it comes to hands-on research. These could include the elimination of manual labor and streamlining the processes of data collection all while possibly introducing bias and reducing accuracy; a double-edged sword. Ganguli et al. (2022: 1754) suggest that LLMs may be the answer to the long-awaited question of how to resolve the large friction present in broad interdisciplinary research and at the same time hold the opposite view, they question the need for interdisciplinary research to assess the impact of the development and application of AI in different social domains. However, they do caution about falling into the prestige pitfall. That of trying to develop cutting-edge technology merely for the sake of perceived institutional advantages to the public eye or monetary incentives without considering the impact that the development and deployment of these models might have on our lives.

Though it could be said that most of the suggestions mentioned so far are theoretical in nature, other authors have tried to stray away from theory and analyze the real-world impacts that AI has had on academia. Livberber & Ayvaz (2023) mention that LLMs such as GPT have already been introduced in many aspects of society including academia. It has already been used in the concoction of academic articles and “information provided by ChatGPT has been used as scientific knowledge” (Ibid). The strong point in Livberber’s et al. paper (2023) is the exposure of academic views, through interviews and surveys, on AI’s such as ChatGPT. These include its anthropomorphic nature, its existence as a positive technological development, and its capacity to shape the future and revolutionize life as we know it. The participants' answers also went in-depth into areas such as its potential as a tool to create significant advantages in the research process and save time, particularly in the literature review and editing stages as mentioned previously (Ibid). While many mentioned AI’s utility as related to brainstorming during the idea generation process, some have expressed concerns regarding possible ethical violations during this same process, only to arrive at the same conclusion previously mentioned by other authors: “There is a need for

legal regulations that clearly define the role of ChatGPT in the production of scientific knowledge...” (Ibid) particularly as it refers to plagiarism, accuracy, and impartiality. And again, we see the same lack of concrete solutions to the call for action.

On the publishing side, Barros et al. (2023) mention the prevalence of inadequate manuscript submissions either due to their undeveloped nature or due to their irrelevancy to the journal to which they have been submitted. They worry that AI tools will only aid this trend of seeking a “quick publish” instead of working towards a high-quality article. But they also mention the possibility of AI aiding in limiting honorary co-authorships granted mainly due to language assistance as a positive impact of AI on publishing trends. This could also lead to a higher integration of the global south in the global north’s knowledge sphere, furthering the efforts of knowledge integration. On a related note, they also mention the use of AI to assist with other academic service tasks such as writing recommendation letters, updating internal policy reports, or preparing specific documentation. Though they do warn about the possibility of rendering administrative and academic staff redundant, a concern echoed by many.

While I’ve attempted to give a separate view of different authors, a lot of the same sentiments and ideas were repeated throughout the different articles. The most salient points would be the characterization of AI as *tools*, the mistaken interpretation of outputs as communication attempts with intent behind them that arise from non-existent NLU capabilities, the elimination of manual labor and the streamlining of data collection, the implications that LMs might have for publishing efforts both positive and negative, a possible breakthrough in the long-standing problem of interdisciplinary research, and the plethora of more overarching generalized benefits related to research (i.e. brainstorming, literature search, content analysis, writing, and editing). All of these topics represent at least a significant enough part of the views of academics on the topic of AI and academia. The sourcing of these, that is, peer-reviewed academic papers, gives them a position that, while sometimes speculative in nature, is nevertheless a legitimizing force. As such, these views will represent the ground upon which a sociotechnical imaginary will be constructed in the following chapter.

2 Speculative Semiotics as a Technique

In this thesis, I will be partaking in a somewhat contentious practice called speculation. I will not engage too deeply with concerns regarding the validity of speculative research, as this is outside the scope of this thesis. I will instead offer a brief overview of some salient works on the matter of the *why* and *how* of speculative research. Additionally, I will offer some insights on the topic of speculation as a technique and not a method. Afterwards, I will briefly go over Mattia Thibault's 2022 paper which delineates a possible way to approach semiotic speculation, which I will employ in this thesis, albeit constrained by the framework of the following section and modified to fit those constraints. Thus, in this chapter, I aim not to defend speculation as related to the academic pursuit. I will merely limit the scope to exposing whose works I draw from as a theoretical framework and the process that I will follow to ensure a systematic and cautious approach to this contentious practice.

2.1 Overview of Speculative Research

“Is speculation a suitable method for social sciences?” is a question posited by Michael Halewood (2017: 52) which is quickly discarded as problematic as it seems to equate speculation with other social sciences *methods*. And a method it certainly is not, as it lacks the rigor and systematic nature of other well-established traditions. However, Halewood warns against discarding speculation's possible contributions by promptly dismissing it, while at the same time warning about using it hastily without much thought (Ibid). Speculation may not be a method but it can indeed have great value as a technique. Wilkie et al. engage with the notion of research techniques and the advantages they can provide to not only bring novelty but to engage with the knowledge construction process in an active manner (2017: 113). They mention two advantages that are characteristic of this approach. The first one is that “... techniques are not burdened with, [...] methodological practices so often presented as ‘self-delimiting self-reproducing or self-explaining’, nor do they necessarily provide a ‘cause, an

effect, or an explanation” (Ibid). While this may seem like a counterpoint to their use, it is intended for the results obtained from speculative research to not inherently relate to the present reality, as they engage with the *becoming* of said reality. The second point presented by Wilkie et al. relates exactly to this matter as it states: “...although research techniques can be bound into configurations of conduct or *dispositifs* [...] that produce subjectification, they evoke the care that is required in developing interventions that may be capable of diagnosing human and non-human novelties into existence” (Ibid: 114). This is such because speculative research has no pretense to deal with nor explain reality in any way, but instead to affect the process of reality construction itself. So while it may be bound to produce subjectification, its intended effect is not to diagnose the present but the potential of its creation.

Savransky proposes that speculation be not taken as an arbitrary use of thought but as a particular and specific practice of thought in the context of an experiment (2017: 26). That is, not a method, but a methodical practice. Savransky conceives speculation as “a wager on an unfinished present, whose potential is that of cultivating thinking to lure experience [...] to take the risk of opening to its own becoming” (Ibid: 26). They mention that for the practice to be needed, a problem needs to present itself as part of an unfinished present that is in the process of developing and becoming, that is, an impending problem (Ibid: 31). As an experiment, taken to be understood as a thought experiment, speculation needs to arise from the possibilities that the facts make apparent through their interrelations (Ibid: 36). In other words, Savransky (2017: 32) expresses that speculation framed as a knowledge pursuit, should not be seen as made-up ramblings of pure imagination, but as extensions of a present that lends itself to possibilities of fact. And most important of all:

“... if speculative experimentations are to become worthy of attention [...] it is not, primarily, because of a truth they might succeed in making manifest, but because of a difference they might be capable of effecting, a mode of becoming felt that might induce a transition into a new empirical situation.”
(Savransky 2017: 34)

This is understood to mean that speculative experimentation, that is, speculation as a technique is not valuable because of its capacity to speak about a state of the present or any truth it may attempt to manifest, but because of its capacity to affect the process of *becoming*

of reality. This concept presented here through Savransky's (2017) paper about speculative pragmatics will be reframed through the lens of speculative semiotics in the following section. But the underlying argumentation remains valid, speculation does not affect the present, nor speaks of a reality to be understood or perceived but interacts with the processes of potentiality of becoming of said present.

Halewood (2017) attempts to approach the subject from a more practical point of view. They present the idea that speculation is not free thought devoid of a metaphorical place to stand upon as "no entity can be conceived in abstraction from the system of the universe" (Ibid: 53). Speculation, they explain, proceeds from "particular actualities with their accompanying possibilities" (Ibid: 54), and these latter ones arise from the relation between ourselves and the world. Successful speculation serves to broaden our understanding of the world and requires that our findings be applicable beyond their initial stage (Ibid: 56). This last assertion stands slightly in conflict with Savransky's position but ultimately could be charitably interpreted to be kindred. If we consider that speculation stands to change the possibilities of becoming of the present, that would indeed entail a broadening of perspective and more likely than not, a fundamental change on the basis of what that present stands to be. What the present is and what it would be upon successful speculation would fundamentally change regardless of whether just the possibilities of the present change or altogether the basis of its own conception. And as much as Savransky emphasized the need for proper speculation to stem from an unfinished present (2017: 26), Halewood believes that speculation demands its own ground, one that cannot be betrayed or transcended and which allows it to exist as an act of thought (2017: 58-61).

Finally, a salient point to take into account is that Halewood mentions three different types of speculation. The first two are named *traditional* and *scientific* both of which have no relevance to the current discussion (2017: 61-62). The third one is *situated* speculation (Ibid: 62). It is called so, as it emerges from a specific starting point that it takes as its ground, which may be perceived as a problem or issue and it recognizes this as its conditions for existence. Halewood calls this the first constraint, which allows for productive speculation (Ibid). The second constraint is that situated speculation involves risk. As mentioned before, this type of speculation recognizes reality as unfinished and in the process of becoming, in

Savransky's words (2017). As such, the risk arises from the fact that the speculation itself is part of that process of "world-in-the-making" and "is inherently and radically pragmatic" (Halewood 2017: 62). And finally the third constraint is that situated speculation has no value judgment, it can not be right nor wrong nor good nor bad, as it is not trying to prove a hypothesis or confirm a reality. Speculation of this kind "is only to be judged in terms of its consequences and effects" (Ibid).

2.2 Speculative Semiotics Next Steps

In a paper published in 2022, Mattia Thibault proposes the use of the term *speculative semiotics* to "...indicate future-oriented speculations marked by a critical approach that is either informed by semiotics as a discipline or has a strong focus on matters related to communication and meaning-making." (Ibid: 3). They cite several precedents for speculative research in general and some for speculative semiotics going back to "The Human Interference Taskforce" of 1891 and some wildly creative speculations about the future of information transmission (Ibid: 1). However, the framing chosen by Thibault results in finding two dimensions to this perspective, one of communicating *with* the future and another of communicating *in* the future, as the main objective of the previously mentioned project was to find signs or texts that would be perdurable and intelligible across long periods of times (Ibid: 4). Neither of these are perspectives that this thesis will be employing. Building upon the framework presented in the preceding section of this chapter, the chosen stance will be that of communication *with* the present *through* the future. Much like Savransky emphasized when speaking about speculative pragmatics, the intention is not to try to engage with an unreachable, potentially non-existent future, but to take advantage of the present and its possibilities of fact, of its becoming, to affect the creation of itself (2017: 26). And in a way, even though never explicitly stated by Thibault, there are hints about this interpretation of speculation in semiotics such as when they say: "The goal would not be to attempt to predict "correctly" the future, but instead to use semiotically-informed speculation to reflect on the present and on the possibilities that emerge from it." (2022: 4-5). Another instance that hints

at this type of perspective on this topic is several mentions of anthropological speculation and its use which, in an almost comical framing, are used to question modern practices of archaeological inquiry (Ibid: 4). In this case, these examples would entail talking *with* the present *through* the past. There is much to be said about communication and the time axis but any deeper inquiry would be beyond the scope of this thesis.

Arriving at a more practical perspective Thibault wishes to not leave speculative semiotics as a theoretical descriptive concept but proposes a series of tentative steps that "... explore the possible synergy between semiotics and speculative research" (2022: 5). The first step is to take a current phenomenon, then "push it" into the future and imagine its evolution, thirdly it is to identify possible techno-social conflicts, and finally crystallize it in a textual artifact (Ibid: 4-5). However given the adjusted perspective suggested above, the steps need to be changed to accommodate a different type of dialectic, that is, one that aims to communicate with the present through the future. My proposition is the following:

- The first step is to establish the current paradigm. To not betray situated speculation as presented by Halewood (2017: 62) there needs to be an established ground from which the speculative practice itself stems, a problem or condition that is perceived as ripe with present potentialities. This requires a careful assessment of the chosen issue or topic that is informed enough to allow for predictions without falling into fantastical thought.
- The second step is to "push it" into the future. As per Thibault's (2022) suggestions, the original scenario needs to be speculated upon, albeit with more constraints than they suggest. Beyond merely being grounded in an 'impending problem' (Savransky 2017: 31), there also needs to be no value judgment (Halewood 2017: 62), and it needs to be a rigorous procedure (Ibid: 58). Halewood warns: "There is an openness involved in speculation so that it does not simply report on what we already consider the world to be." (Ibid). So this pushing of the present needs to be an organic extension of the ground from which it originates. All while maintaining a balance between being faithful to that ground and avoiding the pitfall of what our subjectivities might see as inevitable. In the case of this thesis, this step shall be performed in the form of the construction of a sociotechnical imaginary.

- The third step involves the identification of conflicts. Thibault mentions that a “... sociosemiotic perspective can help look into the deep conflicts and contradictions within future societies” (2022: 4). Be it through sociosemiotics or a different sub-discipline, a modeling of the proposed future scenario should be undertaken as organically as possible. For this, the use of comprehensive pre-established models is desired so as to avoid involving subjective value judgments.
- The fourth step involves bringing the communication back and closing the loop. While it could be done through crystallization via textual artifact as proposed by Thibault (2022: 4) this step could also be performed through a conscious, transparent discussion of the salient elements found through step three. Through this process the reader can travel to this parallel world, experience it, and bring back their own interpretations about the possibilities of the present and as Savransky put it, “a mode of becoming felt that might induce a transition into a new empirical situation.” (Savransky 2017: 34).

Based on this adjusted system the steps shall be applied to this thesis in the following manner. The first step, establishing the ground, has been completed through the outline presented in chapter one of this thesis. Step two, the construction of the sociotechnical imaginary, will be laid out in the following section of this chapter. Step three, identification of conflicts, will be done in two parts. In chapter three the model to be used will be presented. In chapter four, the model in question will be used to analyze the sociotechnical imaginary presented in chapter two. This will entail using the ground established in chapter one to find the basis of the analysis, the issues identified that serve as stems from which the sociotechnical imaginary was constructed in chapter two. Then applying the model presented in chapter three to break down the issues into parts, that is, the three moments of the dialectic of the social construction of reality. Lastly, we will compare the moments of the dialectic as they stand in the present academic processes as opposed to as they would be in the sociotechnical imaginary. By following these steps we will be able to identify divergent patterns of knowledge construction that arise from the implementation of AI into academic processes. Step four, crystallization, will be performed through the discussion section at the end.

2.3 Speculating on AI and Academia

Based on the views presented in chapter one about the current perceived state of AI, both from specialists and from the academic community at large, we can take that to be the ground from which this situated speculation would stem. This section will thus constitute the proper construction of the sociotechnical imaginary which will be broken down for analysis in chapter four. As mentioned previously this thesis will assume that all instances of LM deployment would be done with no outright malicious intent, as that would be counterproductive to analyzing the best-case scenario for the potential involvement of AI in academic processes and research. However, this does not mean that actors do not have different incentives, that they strive for different outcomes, and that some of these efforts would be detrimental to other actors. A good example of this is the complaint from specialized academics about the power of the private sector in AI development and how their incentives do not align with what they perceive to be in the best interest of the public (Ganguli et al. 2022: 1756). And while some mitigating actions might be taken, these asymmetries in incentives are nigh impossible to completely overcome. This type of mismatch would not be considered malicious intent but merely different parties looking out for their own self-interest. This will not be central to the speculation efforts but nevertheless is important to note. In the following paragraphs, the different axes of possibilities of AI development will be explored.

The first point to be analyzed is that regarding the distribution of power. At the moment the barrier to entry into the development of LLMs is increasing constantly due to the amount of data needed for the training process (Bender et al. 2021: 619). This has led to very few models dominating the market and concentrating the power of decision-making only on a few entities. The private sector, which is incentivized mostly by profit, is the main force behind the development of LM technology. This has certainly skewed its initial growth towards the desires of investors who are constituted primarily by the industrial sector. As this technology grew, it attracted the attention of different stakeholders sparking interest all throughout different sectors of society, some of which are calling for governmental

regulations to some of the more salient issues (Ganguli et al. 2022: 1757) even though the legislative force has been slow to act upon it. We can speculate that this trend will continue and the market of LLMs will be dominated by few alternatives controlled by elite individuals with few governmental controls and regulations. However, as LLMs continue to spread to our everyday lives these private entities will be forced to take into account the interests of other stakeholders besides the industrial sector. So while the power may be centralized and control may be minimal the widespread nature of this technology may just lead to an apt representation of the interests of the public in its development and maintenance. We can see some examples of this already occurring as people are calling for these companies to solve issues of bias and prejudicial views currently present in the models (Bender et al. 2021).

On the topic of knowledge scope, in chapter one it was established that the current trend for AI is that of large language models that are of the generative and transformer variety. GPT being the current trendsetter. There are a lot of efforts being directed toward the development of generalized conversational models (AlZu'bi et al. 2023). Koubaa suggests that one of the possible axes of future development would be that of “domain specific language training,” insisting that models trained on domain or industry-specific data could be produced to better suit each one, healthcare, finance, or law being the salient examples (2023). However, immediately after that, Koubaa et al. themselves suggest that this technology also has the capacity to be a model that can transfer knowledge between domains (Ibid). Considering this, and the current trend of conversational AI development, it is likely that the incentives of the private sector, who are developing the models themselves, are more aligned with the creation of an all-encompassing model of conversational AI that is versatile enough to be able to be marketed to all sectors of the population. This AI would not only be capable of communicating in a conversational manner to fulfill the general public's expectations created by the promises of the industry itself but also have the potential to produce industry-specific language modeling when prompted. This outcome would align the incentives of the different stakeholders making it the most likely outcome. At the same time, it should be recognized that this would probably mean that the AI would need specific prompting to produce certain specialized results and it would represent a barrier to some sectors of the population. Nevertheless, in regards to the scope of knowledge, a large

centralized single model with the capacity to shift from different registers when commanded seems like the most likely outcome given the current trends and the incentives of different parties. This could have an effect in particular on how we understand interdisciplinary research.

One of the less discussed axes of analysis is that of language. This is by no means a minor issue. AI models have been trained on different natural languages with varying degrees of success. More dominant languages, those that have the most active users and have thus produced more data, have an advantage when it comes to training language models. This is not only regarding the issues of inter-lingual translation as evidenced by some real-world wrongful arrests caused by these inaccuracies (Bender et al. 2021), but also an issue of representation and accessibility (Koubaa et al. 2023). Some suggestions on how to improve upon the modeling of so-called “low-resource languages” have been brought forward, but the confidence projected in those solutions seems low as the possibilities explored look to be tentative at best as showcased by the emphasis on the necessity of developing “new” training techniques or evaluation metrics and benchmarks (Bender et al. 2021: 616; Koubaa et al. 2023). If we were to take this current situation as the ground, it would follow that while some development in this area could take place it is more likely that this bridge might not get completely overcome. It is probable that high-resource languages such as English will continue to dominate the market by showcasing the best results for LLMs and that other low-resource languages while supported will not show comparable results. This would be despite any stakeholder's incentives for the opposite outcomes as for the moment it seems to be a limitation of the path of technological development chosen, that is by scale. If more efforts were put into the development of these uncertain “new” methods to try to lessen the disadvantages of low-resource languages maybe this would not be an issue, but at the moment there is little incentive for the private sector to put emphasis on these languages that are more likely spoken by populations who are already poorly represented and hold little power. This has the potential to lead to asymmetries in access and representation that are present today and while they might lessen with enough effort, even then they would probably not disappear. But we should not discount that even if it is not perfectly symmetrical, this technology could increase the accessibility of certain populations to information and even the reach of their

voices beyond what could have been without it, both outside the academic world and inside it as mentioned by Barros et al. (2023).

This asymmetry as presented will also be affected by another axis of possibility, that of the competency of the users. Several researchers have noted the increased importance of prompting skills and the possibility of this developing into a full-fledged area of expertise (Livberber & Ayvaz 2023, Koubaa et al. 2023). This is a highly likely scenario that mirrors the increased importance of search engine use proficiency after the advent of that particular technological development. This would also suggest that prompting might become a sought-after ability that might eventually even develop into an outright specialty or area of expertise. This would run counter to the previously mentioned increase in accessibility, making its casual use widespread while its specialized use less so.

One, if not the main concern listed in chapter one by both specialists and non-specialists alike was the increasing prevalence of biases on the outputs of the language models. This was brought upon by increased scale and by the same mechanisms that allow some of the greatest advancements in this technology. And while many have suggested that maybe scaling models in this manner is not the only way to advance (Ganguli et al. 2022) the current direction of development seems to indicate the opposite. And given the increased prevalence of biases that threaten not only to re-victimize individuals but to reinforce said biases in the population there have been several suggestions about how to deal with this issue. Some of these solutions seem to be technically feasible (Bender et al. 2021: 615) so there is no reason to doubt that this issue might be at the very least successfully mitigated. However, human biases and prejudices are, historically speaking, quite durable as they are a matter of social construction and maintenance. So in regards to this particular issue, the only possible solution is an ongoing battle of reinterpretation and sanitization of the models. And given that this issue is one of the most commonly referenced there is no reason to doubt that proper attention and effort will be directed to this area of development. As such, it would be an appropriate prediction to say that the onset of deepening biases on LLMs will be mitigated and treated as an ongoing struggle with this particular type of technology because of its nature and quality of imitating human speech which is inherently full of of biases. If continued efforts are put forward to mitigate the effect of this emergent risk to minimal levels there is no

reason to believe that the effect would be different to that of being exposed to any human-created text.

The second issue most echoed by the texts referenced in chapter one is that of a deep concern for plagiarism. While this thesis does not deal with plagiarism concerns as related to training data used, the question regarding the use and apprehension of the outputs of the AI is an entirely different issue. As mentioned before, the outputs produced by LMs are not a result of any entity that can have accountability for what is produced (Bender et al. 2021: 618). While we may attribute emergent abilities to LLMs that are akin to reasoning (Wei et al. 2022) we have to recognize that these models, while they may be capable of natural language processing, are not capable of natural language understanding (Bender et al. 2021: 615). That is, while it may look to us like what they are doing is just like thinking and rational thought, it inherently is not. We may treat it as such since the phrase goes: “If it looks like a duck, walks like a duck, and quacks like a duck, then it just may be a duck.” While some might know it is indeed not a duck, and as it stands, LLMs are not capable of NLU, by any metric of human perception it might not matter. We are wired to attribute meaning to language and intent to entities that are responsible for it (Bender et al. 2021) and we might try to fight that instinct but it will be an uphill battle. At the end of the day, the speculation regarding this particular axis is deeply related to how we interact with AI and what role we assign to LLMs as a society. This particular topic will come back in the following chapters with increased relevancy because it seems that this is the core of the current discussion regarding LMs. What is the role that they should have? As Karen Spärck Jones wondered already in 2004, we may accept it as a noisy channel, disregard the question and relegate its existence to *just a tool*, or ignore the lack of correspondence between outputs and meanings, and it is hard to find a fourth option. Either way, the question regarding the legitimacy and epistemological backing of the outputs of LLMs is what ultimately is still being debated and the hardest of all the axis of possibility to speculate upon. I will adventure to say that given the number of references to AI as a tool and the types of uses it is speculated it could have, we seem to be going in the direction of relegating the existence of AI to “just some convenient technology” (Spärck Jones 2004), a tool for our convenience and to which we assign little afterthought.

All of this would lead our speculation of the future of AI to be:

- An intensification of the development of generalized conversational models which would have the capacity of code-switching between different domains and potentially translating between them.
- A somewhat symmetrical representation of the interests of different stakeholders as the ubiquity of the technology increases all throughout society.
- An asymmetrical representation of different languages based on the amount of available resources for training, leaving poorly represented populations that speak low-resource languages in an unfavorable position.
- While this technology would lead to an increase in accessibility to information and language assistance from less advantaged populations, it would maintain its highest performance firmly hidden behind prompting expertise and AI competency.
- An ongoing struggle to manage biases and prejudices in the AI's outputs, not dissimilar to human efforts to do the same at societal levels.
- The role assigned to AI seems to be leaning towards regarding it as a tool for our convenience while the epistemological validity of its outputs remains a matter of contention.

3 Knowledge as a Social Endeavor

Berger and Luckmann's treatise on the sociology of knowledge published in 1966 under the name "The Social Construction of Reality" will serve as the framework for the analysis of this thesis. They portray the social construction of reality as a dialectic process with three moments: externalization, objectivation, and internalization (Ibid: 70), all of which rely on language as a key component that aids the different moments of the dialectic, albeit in distinct capacities (Berger, Luckmann 1966). The authors emphasize that while the dialectic of social construction is composed of these three moments, they are not organized in a temporal sequence and cannot be understood separately (Ibid: 149). Berger and Luckmann intertwine these three moments with the processes of institutionalization, legitimation, and socialization, among others lesser ones. But for the purposes of this thesis, the overarching three moments of the dialectic will be preferred, even if the other processes are mentioned and utilized occasionally. This choice has the intention of isolating the semiotic processes at the individual level while still being able to extrapolate their social-level counterparts. At the same time, because the moments of the dialectic are so intertwined, it will be difficult to speak about them in any type of sequence without mentioning elements of each other and ultimately how they all interact with language. In this chapter, I will briefly explain the workings of each of the three moments of the dialectic while mentioning the way in which they interact with one another. Later on in this same chapter, I will delve deeper into language as a key component of this process.

While not developed in this chapter, other key concepts will be needed and employed in the analysis of the following chapter. These concepts do not have the intention of providing additional methodological tools but to aid the concepts presented in this chapter and fill some gaps that were left undeveloped by Berger and Luckmann. In other words, these additional concepts are meant to complement and not substitute anything presented in this chapter. These include:

- Subjectivity in language: as developed by Benveniste (1971) to speak about the subject positioning inherent in language processes and aid in the understanding of the process of crystallization of subjectivities as explained in section 3.4.
- Inseparability of form and content: as developed by Voloshinov (1973) to bridge the gap between the form of a sign and its ideological signification. This will allow the analysis to engage with LLMs outputs as having epistemological weight even if by proxy of the form they are taking and as such, to partake in the dialectic of the social construction of reality.
- What is an author and the author function: as developed by Foucault (1998) to bridge the gap and further develop the idea of how LLMs outputs could be considered as fulfilling the author function and where they lack ground to stand upon. This last author will be used in the last subsection of the analysis chapter which presents different concerns than those of the previous sections.

3.1 Externalization

Berger and Luckmann present externalization as an anthropological necessity of the human being (1966: 70). They say that because of the natural internal instability of the organism, human beings feel the necessity to provide a stable external environment to inhabit, and thus a social order stems from biological reality, even if it is not by direct causality (Ibid). Externalization is deeply ingrained into the process of institutionalization. The first instances of institutionalization arise from the process of habitualization undertaken by human beings with the objective, conscious or unconscious, of providing relief when faced with constant decision-making processes (Ibid: 71). Specifically, “institutionalization occurs whenever there is a reciprocal typification of habitualized actions by types of actors” (Ibid: 72). All of this to say that externalization is the consequence of the necessity of humans to order the world, control nature, and organize chaos. This necessity, while not caused by biology, is influenced by it and inevitably humans as they exist will start to find patterns in the world around them. These patterns will eventually be turned into structured forms of thought that, when shared

with others, give rise to the foundations of social order. This is roughly how externalization contributes to the construction of social reality.

As mentioned before, the three moments of the dialectic are not temporally sequential but nevertheless intertwined and significant. As it will become clear later in this thesis, AI and the types of interactions that humans have with AI that are pertinent to this analysis, that is, user interactions, are more intricately woven with the other two moments of the dialectic. But externalization is nevertheless crucial, as it pertains to the projection of an individual's reality onto the world in the conception of symbolic universes (Berger, Luckmann 1966: 122), that is, the projection of internalized meaning into reality. At the same time, knowledge programs the channels through which externalization processes construct reality through objectivation (Ibid: 84). So even if externalization is not the focal point of this analysis, it is nevertheless an important component to keep in mind when utilizing the other two moments of the dialectic.

3.2 Internalization

Internalization is posited as related to socialization. Both primary and secondary socialization processes are presented in the text as part of the internalization moment of the dialectic (Berger, Luckmann 1966: 149-150). Primary socialization occurs when an individual has internalized the objective world of a given society to the level of taking over said world as their own, and is the most important one, as it sets the basis for all future possible socialization which will have to conform to it. (Ibid: 149-151). Secondary socialization, in turn, “is any subsequent process that inducts an already socialized individual into new sectors of the objective world of his society” (Ibid: 150). As such, primary socialization is more threatened than secondary socialization in marginal situations (Ibid: 168). In this case, secondary socialization will be a more significant process to the analysis as any form of specialized knowledge is not pertinent to primary socialization. This latter one is only concerned with the induction of the individual to the core of the symbolic universe while academia would be certainly aligned with sub-universes that have their own institutional processes of reality maintenance. Nevertheless, both processes require identification.

Internalization itself occurs as identification occurs. It is the result of a “dialectic between identification by others and self-identification, between objectively assigned and subjectively appropriated identity” (Berger, Luckmann 1966: 152). In other words, internalization requires a social order that is pre-established and already objectivated and from which the individual can draw upon through their significant others. In this way, internalization is related to objectivation as the process of identification that occurs through significant others and draws from the socially constructed reality, needs a previously objectified world in general and institutions in particular. This internalized reality is maintained through externalization processes, as projecting internalized realities and validating them through processes such as conversation are everyday reality maintenance mechanisms (Ibid: 172).

3.3 Objectivation

Objectivation at a simple glance can seem like the most significant of the three moments of the dialectic but it could never exist without the other two. Objectivation consists of the process of ordering the world into objects that can then be apprehended as reality (Berger, Luckmann 1966: 84). Objectivation is the mechanism through which the social stock of knowledge of society is created. But by itself, it cannot functionally exist, as without externalization there would be no projections of internal subjectivities which to build bridges upon to create an intersubjective reality which can then be perceived and apprehended through internalization. If taken to the extreme this leads to reification which implies “that a man is capable of forgetting his own authorship of the human world [...] and that the dialectic between product and producer is lost to consciousness” (Ibid: 106). In other words, reification is the culmination of an objectivation process to the point where humans forget that they are the creators of a particular product or reality and proceed to experience it as a facticity independent of their own creation. While not extensively treated in the text, reification will be treated as a significant possibility concerning AI in the analysis section.

Legitimation as a process is a second-order objectivation of meaning (Berger, Luckmann 1966: 110). The general aim of this process is the integration and consolidation of the totality into a whole (Ibid) in particular in regard to institutions. The ultimate form of this type of process is the creation of a symbolic universe that justifies and provides a holistic understanding of reality to the social group (Ibid: 114). This symbolic universe, in modern societies, is conformed of a plurality of partial universes that coexist in mutual accommodation with a shared core universe (Ibid: 142). For groups that have specialized knowledge and which partake in secondary socialization practices, sub-universes can be formed (Ibid: 113). As mentioned before these sub-universes are considerably less threatened by marginal forces as they are acquired through secondary socialization processes leading to a lower facticity level. Nevertheless, these sub-universes are those that conform to academic domains and practices and are of key importance to this thesis.

Language plays a big role in objectivation, almost becoming critically significant to the moment itself. It is through linguistic designations that have the capacity to abstract experiences from individuals and their biographical occurrences and, in the process, anonymizing them, that objectification becomes factually feasible at the scale that is employed in modern societies (Berger, Luckmann 1966: 86). And in turn, language also allows for the objectification of new experiences allowing their quick incorporation into the social stock of knowledge (Ibid). But truly language plays a much larger role in the dialectic. While appearing to not be indispensable, it would be borderline impossible to establish the type of institutions that humans have in modern societies without the use of language.

3.4 Language and the Construction of Reality

Language is a crucial component of the triadic dialectic of the social construction of reality. It plays a role in every moment of the dialectic but also in the mechanisms that serve to establish, maintain, and also to modify it (Berger, Luckmann 1966: 172). This is primarily due to linguistic signification. Signification is a particular instance of objectivation which consists of the human production of signs that can be differentiated from other objectivations

by their “explicit intention to serve as indexes of subjective meaning” (Ibid: 50). Even then “signs and sign systems are objectivations in the sense of being objectively available beyond the expression of subjective intentions” (Ibid: 51), that is, they are detachable and thus, common objectivations are maintained primarily through linguistic significations. Language goes on to play a multitude of roles: a repository of experience, a method to crystallize subjectivations, a coercive force of objectivity, a way to transcend beyond space and time (Ibid: 50-54), etc. But in whichever way, language is undoubtedly intrinsically woven into the dialectic process of the social construction of reality as “the edifice of legitimations is built upon language and uses language as its principal instrumentality” (Ibid: 82). This is because “language provides the fundamental superimposition of logic on the objectivated social world” (Ibid). In other words, language is a coercive force that crystallizes subjectivity and makes it conform to an intersubjective reality which sediments itself into an objectivated reality that will go on to create a symbolic universe. Language frames most of the processes of our social interactions and bridges the gaps not only between each other but between ourselves at different times and places. It provides a framework for us to perceive and categorize our socialization practices and allows us to apprehend the world in an organized manner. Language is also the main method of reality maintenance as it is through conversation that we realize the world by actualizing our inner subjectivities while at the same time projecting them.

But language not only serves to become the deposit of large aggregates of collective sedimentation (Berger, Luckmann 1966: 87). Language can also serve to objectivate specific shared experiences that are available to a few and create linguistic communities (Ibid: 95). These linguistic communities can participate in secondary socialization practices that require the acquisition of role-specific vocabularies (Ibid: 158). These roles appear as institutional representatives and mediate between the institutions and the aggregates of knowledge, as well as serve specific purposes of not only initiation through secondary socialization but also reality maintenance, as these sub-universes are less stable and are subject to competition (Ibid: 91-103). As mentioned before it is these types of linguistic communities and their associated sub-universes of meaning that conform to different areas of the the academic reality. Taking advantage of the inherent plurality of modern social configuration mentioned

before we can extrapolate that the same type of configuration can be used for academic knowledge. While sharing a core common symbolic universe, the different domains of knowledge, disciplines, communities, groups, etc. partake in their own smaller sub-universes of meaning. This gives rise to a plural landscape within the overarching stable and organized academic reality. This is not to say that academic knowledge stagnates due to its stability but instead that it endures and remains steadfast in the face of its inherent internal plurality due to a stable core symbolic universe.

4 Analysis of a Becoming of the Present

In chapter one, several salient topics were identified as possible sources of issues in relation to the implementation of AI into academic processes. These issues will now inform the sections of this chapter to ensure that the analysis remains focused. At the same time, these issues also served as the ground from which the sociotechnical imaginary was constructed in chapter 2.3. As such, this constructed imaginary future will frame and inform the analysis of this section. The process will be as follows. In each section of this chapter, a different issue or category of issue will be presented and developed based on the sociotechnical imaginary constructed. That is, how would a certain problematic area exist in this potential future? Then, the moments of the dialectic will be identified for each, the present processes, and the potential future processes. Lastly, these two will be compared in themselves and between themselves to identify deviations in the processes of the social construction of knowledge in the potential future as contrasted to the present.

As such, the different issues identified in chapter one give name to the four sections of this chapter. In this way, each section represents the limits of the imaginary plane as per the chosen ground. In summary, this chapter aims to analyze how the moments of the dialectic and the processes of knowledge construction as presented in chapter three change between a future in which the processes of academia remain as they currently are, and those as conceived in the sociotechnical imaginary concocted in chapter two using chapter one as the ground. For this purpose, the issues identified as the ground and summarized in the last paragraph of chapter one will be divided into: fetching data, writing and editing, interdisciplinary research efforts, interlingual knowledge creation, and epistemological concerns of outputs.

4.1 On the Possibilities of Fetching Data

While fetching data is not mentioned in this form anywhere in chapter one, concerns regarding data collection and literature review are mentioned repeatedly in a positive light. I have put these two potential issues under one category of issue as they both involve similar processes. As such, fetching data will be considered to encompass all processes of acquiring different types of data for research purposes. These include those mentioned in chapter one and recognized by specialized and non-specialized researchers alike as possible areas where issues might arise, namely: literature search, literature review, data collection without manual labor, data collection through manual labor, and additionally, content analysis. This latter one is understood in the literature as the analysis of the content of the retrieved data and not as the analysis of content produced by the user, that is, as an extra step of the data acquisition process that is inevitably attached to the outputs that LLMs provide to these types of queries. However, even if processes such as literature review and data collection involve the retrieval of data used for research, the manner in which they function is different in practice and a distinction needs to be made. Because we are dealing with the process of knowledge construction in society the distinction is better placed at where the data comes from. That is, does it come from a source where it has been previously processed and objectivated or is it “new” data, freshly collected, and which has not been objectivated yet? As such, for the purposes of this analysis, we can make the following distinction in data fetching processes. Data acquisition would entail literature search, review, and analysis, that is, fetching data mostly from texts where it has already been objectivated and is presented through symbolic means. While data collection would involve the process of collecting “new” data either through manual labor, that is, experimentation and sample collection, or other types of empirical means, that is, data that has not been previously objectivated.

Data collection is presented in the papers in chapter one as a promising area for the use of AI, particularly by Koubaa et al. 2023. The main difference to be found as compared with data acquisition, is that data collection would necessitate the creation of specific AI models that may not even involve NLP at all. This is because, depending on the type of data

and the type of experiment or procedure needed for its retrieval, different AI models might need to be implemented. As it stands, these types of procedures are codified and performed methodically mostly by humans but also by machinery depending on the specific domain of knowledge. It can be assumed that in our sociotechnical imaginary, the individual researcher or group of researchers when confronted with a market in which the main technology is a generalized conversational model as assumed by the speculation in chapter two, would need to find a solution that is an alternative to the mainstream. As such, these imaginary researchers would need to find the closest appropriate AI that can perform the needed tasks or develop alternatives that are suited for the tasks and that mirror the already established procedures in each domain of knowledge. Given that researchers in the present time already are using advanced technology for this type of data analysis and collection, the potential involvement of AI that is suited and tailored to individual tasks would ensure that the procedures are as close to the current ones as possible while maximizing efficiency. In other words, generally speaking, the involvement of AI in data collection would not change the knowledge creation processes in any significant way because the type of tasks necessitate the creation of specialized models that closely resemble pre-established practices ensuring that the methods and their reliability remain as unchanged as possible.

On the other hand, data acquisition that takes the form of literature search, review, and analysis is quite different. This type of procedure could very easily be performed by the type of generalized conversational model speculated to be dominating the market in our sociotechnical imaginary. The prototypical process of data acquisition involves a researcher who already has a query and a general direction and is looking for material that is relevant to their research theme or topic. As it currently stands, the researcher has a number of possible avenues of discovery, whose availability depends on the particular sociocultural background of the person. The large majority of them would involve one of four widely known options: internet search, library search, interpersonal relationships, and citation tracking. The last two however are not the focus of this analysis as they are assumed to remain a constant in our sociotechnical imaginary unlike the first two which are the ones that could potentially be supplemented or entirely replaced by LLMs. Sources born out of interpersonal relationships need little to no explanation as they encompass any literature found through the

recommendation of people the researcher considers reliable. Citation tracking, sometimes also called citation chasing, would involve utilizing an already known relevant text and its bibliography to find other relevant texts, a process known as backward citation searching, or looking up other texts that have cited said text, known as forward citation searching. Both of these would not be under threat by the introduction of LLMs for data acquisition. On the other hand, we have internet search, that while not currently utilizing LLM technology, already involves a different type of AI technology as mentioned by various authors, among them Livberber & Ayvaz 2023 and Koubaa et al. 2023. Most search engines use some type of algorithm or similar technology to fetch relevant data based on input. This might not seem to be significantly different to utilizing LLMs but the key difference is the way in which information is processed in order to identify relevant data and the form of the output. Most search engines rely on keyword similarity for the bulk of the search queries and retrieve the data as a link to the source as is (Google 2025). This process could be considered similar to that of a library search in which the individual uses the general domain of knowledge, topic, or other keywords to find relevant literature from that which is available at a particular repository. Both of these methods of data acquisition start from a theme, query, or otherwise known general topic or direction and attempt to retrieve relevant literature from a particular pool of texts. This is the type of data retrieval that would be most similar and possibly replaced by the introduction of LLMs into the process.

Data acquisition could be modeled as a process that starts with an externalization moment, in the form of the selection of the theme, topic, or area of study. This moment would involve either the selection of the closest match from those provided by a library, the formulation of a short sentence or collection of words that would serve as the input of the search engine, or the composition of a short query to input into the LLM. At present the conditions that enable the selection of a particular library or search engine will be taken to be analogous to the one that leads the individual to the selection of a particular LLM, i.e. accessibility, familiarity, availability, etc. After the initial moment of externalization, the selected method will result in the retrieval of numerous literature options from which the individual must choose. The results of each method might vary in form, from literal books and papers that the individual has to sift through, to a large list of links that lead to digital texts or

paywalled content, or even a text that offers an orderly list of relevant sources and possibly an explanation of why they are relevant and where to find them. But what follows in the dialectic of the social construction of reality is the same. The individual must decide whether the fetched data is relevant to the externalized objectivation by comparing the results to the expectations they hold based on the reality they inhabit. In this moment of the dialectic internalization occurs as the researcher evaluates the literature presented through the different methods by engaging with the particular sub-universe that they inhabit as a result of the secondary socialization practices that lead them to acquire the vocabulary and objectivations from which the original externalization stemmed. There is no difference in the processes of knowledge creation and the moments of the dialectic between an internet search and the utilization of an LLM for literature search, as they both rely on a language-based input to provide a language-based output. Without the initial externalization provided by the individual, there would not be an output to which to apply an internalization moment. The output of the LLM is created directly from the input as an extension of the language choices made by the individual. This is evident, as explained in chapter one, given the nature of LMs and particularly transformers as they function by predicting each word based on the previous N words (Bender et al. 2021: 616). Thus, the starting point of an AI process query would involve the N words inputted by users which based on the algorithm and the data set used for training would yield different results. Nevertheless, there is a causality between the input and the output that is undeniable. As such, the language choices made by users, as explained in chapter three, are objectivations that are detachable and serve to bridge a gap between being the expressions of subjective meaning and a coercive force of objectivity. At the moment of externalization, the words are working as expressions of subjectivity, and at the moment of internalization, the same words are a reality-maintenance mechanism that we apprehend and through which the legitimization of the reality we inhabit occurs. Thus, whether the words we receive are those of an LLM output or the ones retrieved through a search engine makes no difference because it is the individual who will perceive those words and attribute to them the corresponding level of facticity according to their own apprehension through the internalization and objectivation moments. While the mediator of the information may change from the internet or the library to an LLM, the processes through which the information is

processed are remarkably similar. As such, there would be little to no repercussions at a societal level, and taking the analysis to that point would result in the same conclusion.

4.2 On the Processes of Writing and Editing

The process of writing and editing involves the creation and or change of a written text. In this regard, it differs greatly from the processes described above as this one deals with the final product of an academic pursuit. It is the crystallization of all the work into a written artifact. Both of these, writing and editing, are also mentioned by the literature presented in chapter one as potential areas where issues might arise, though the opinions were less optimistic than those regarding data fetching. While most authors regarded both of these areas positively, there were some dissenters that warned of possible issues. But at the same time, the sociotechnical imaginary constructed previously would easily accommodate this practice. The generalized conversational model that is speculated to be ubiquitous in this imaginary future would lend itself with ease to intervening in writing and editing processes as both of them deal with language at their core, a strong suit of such an AI model, as explained in the literature outlined in chapter one.

The process of writing and editing without AI does not need to be explained as it is self-explanatory. The only caveat would be to highlight that certain technologies like “autocorrect” and “grammar checking,” as done through word processors like Word from Office, have been ingrained into writing through typing for quite some time and these technologies have evolved to sometimes include AI features (Schaul 2024). So we are using AI to write more often than not, even if we don’t realize it. The process of writing and editing with AI, however, can have several different levels of involvement as mentioned by Pinzolit 2023. From spellchecking, proofreading, syntax editing, structural editing, or outright writing or rewriting whole texts, the extent of AI penetration in the process varies widely. For the purposes of this section, we will separate them into three categories. Minor intervention, for processes such as spellchecking and proofreading. High intervention will be reserved for cases in which the AI will outright write or re-write whole sections of text. Medium

intervention will be used for cases such as semantic editing or similar processes that go beyond spellchecking but not so far as to have the AI write the text. These latter ones could also involve feeding a particular written text to an LLM and asking it for suggestions or improvements to the text and proceeding to apply them.

As said before, the process of writing and editing without AI needs no explanation as it is extremely straightforward, even if endangered nowadays. Translated into the model of the social construction of reality presented in chapter three and taking into account just the writing and editing and not any step that led to this process, we could say that the main moment of the dialectic employed is externalization. After all the other research steps are completed, writing involves gathering the processed information and externalizing it through linguistic signification, which is in itself an objectivation. Editing on the other hand is the process of cross-checking the actualized language choices with the expected ones, as such, it also involves the moment of externalization, albeit tinted by an internalization to a lesser degree. In both of these cases, the externalization would be informed based on a larger universe of meaning and informed by a particular sub-universe of meaning. This latter one is inhabited by the individual, and they have gained access to it through secondary socialization processes. That is, when writing or editing by themselves, the individual will externalize the language according to the rules and directives previously internalized when they gained access to a particular linguistic community, in the case of academia, the larger community would be a language-speaking community (e.g. English speakers), and the smaller community or communities acquired through secondary socialization could be a particular field or area of study, or the requirements of a particular institution or journal.

When AI is involved, as mentioned previously, it could be in one of three levels of penetration. This is significant as any type of intervention in the externalization moment of the dialectic, particularly in the processes of writing or editing, would be changes in the *form* of the resulting text. And because AI is regarded as dealing with form and not with content (Bender et al. 2021), in this way, one could mistakenly conclude that any intervention at the level of form would hold no impact on the content of the text and would thus not affect the academic validity of the ideas expressed in said text. But form and content are inextricably linked. As explained by Voloshinov: “The theme of an ideological sign and the form of an

ideological sign are inextricably bound together and are separable only in the abstract.” (1973: 22). We should thus consider any involvement of AI in the construction or modification of a text as affecting not only its form but also its content (through its form). However, there is a certain margin that we could potentially consider acceptable, as it is hard to argue that, for example, correcting a typo would constitute a significant enough change to affect the validity of the content of an academic text. As such we will proceed to analyze the different levels of intervention mentioned at the start of this section.

In the case of minor intervention, the AI would simply proofread the text. This could result in modifications such as the correction of spelling, punctuation, grammar, or other similar changes. At this level the penetration of the AI in the process is minimal, and the changes made are being actualized in comparison to linguistic rules made at a very wide social level. That is, any change made in regards to spelling, punctuation, grammar, or similar would be a change that any member of the same language-speaking community would agree with, as opposed to members inhabiting more niche sub-universes of meaning. This is, of course, a generalization, as dialects and other variations in language exist, but in general for communication to occur two members of the same language-speaking community need to agree on a set of rules that dictate how the language works. It could be argued that this type of minor intervention has minimal impact on the content portrayed, as it does not deal with sentence structure or semantic changes. Considering that sub-universes of meaning are more threatened by marginal situations and as such, their significations are not as stable as those of the core universe, that is, the main language-speaking community, semantic changes would represent a much larger threat to objectifications than spelling or grammar. So, while there’s still an underlying danger as any change in form is a change in content, this level of AI involvement is largely innocuous.

When it comes to the level of high intervention the case is the opposite. By allowing the AI to outright write or re-write whole texts it is directly intervening in the externalization process and having maximum impact on the form and the content of the externalization. However, it would be difficult to call this type of interaction with the AI an externalization moment in the same way as the one previously described. In this case, the process could involve an initial externalization required for the composition of the query that serves as an

input for the AI, followed by a secondary externalization moment, more akin to that of editing than writing, in which the researcher would cross-check the output of the AI with their internalized expectations. As such, this type of AI intervention results in the researcher in the position of an editor for a text that is either authored by the AI or without an author, more on this in section five of this chapter. Either way, it would be difficult to argue that such a text could be presented by the researcher as their own without endangering their own ethical standing befitting their position. As such, it could be said that this means this type of intervention would not be practiced in our sociotechnical imaginary. But that would mean underestimating the power of normalization. AI is already being used to write everyday texts such as emails, notifications, complaints, etc. However, it has been suggested that AI would be a potential tool to aid article writing by asking the AI to produce paragraphs that would later be assembled by a human into a text (Livberber, Ayvaz 2023). As these technologies encroach into everyday activities they start to become normalized and we start to depend upon them. And there are not many steps between writing paragraphs to be assembled and just writing the text. And in a world such as that described by the sociotechnical imaginary where a generalized conversational model has become a ubiquity in our everyday life processes, it would probably feel natural to use it for just one more thing. So it is important to not discard possibilities that might seem outrageous given the present context and to evaluate the whole range of possibilities.

Between minor intervention and high intervention, there is a lot of middle ground. I've called this middle ground, medium intervention, and described it mostly as structural editing and semantic changes. But in reality, there are a lot more variations and possibilities that stand between just spellchecking and outright writing the text. It is difficult to ascertain the entire range of possibilities and what these might, in turn, entail for the writing process. So what we can say is limited to knowing that these processes would stand between the ones that come as a result of minor interventions and high interventions. That is, they stand between being innocuous and being of great ethical concern. However, in our sociotechnical imaginary, we had established that there would be a generalized conversational model that would be ubiquitous throughout a particular society. As such, the types of queries that could be asked of the AI would extend the whole range of interventions but would most likely lie somewhere in

between rather than at the extremes. Particularly if institutional regulations were to be implemented. This is a case in which the use of AI would probably end up being regulated by institutions in the form of allowed and non-allowed queries that can be made. As the institutions themselves would have a level of self-preservation that would lead them to try to control any practice that endangers their own standing. An example of this could be a particular university allowing inputs such as “Please correct punctuation and spelling in the following text.” and disallowing more ambiguous queries such as “Please correct this text.” The resulting impact of AI in writing and editing processes would be differentiated by which process it is being applied to and to which extent. If high interventions could be considered as renouncing the authorship of a particular text to the machine, so to speak, then those could be and probably would be disallowed by most institutions in our sociotechnical imaginary. But interventions that would lie somewhere between minor and medium, and could be categorized as editing, would probably be allowed through regulatory practices, in particular in the case of texts for publication, as the advantages of these technologies would be weighed against the dangers they pose. All of this would take place assuming that the relevant entities would be concerned with the authorship rights of academic papers and with plagiarism, both of which have been mentioned as current concerns in the literature presented in chapter one.

One important caveat to this whole section is the reminder that in our sociotechnical imaginary, there is an asymmetrical representation of different languages. I have made several mentions of language-speaking communities but it would not be the case that all language-speaking communities would yield the same results from this type of interventions. Low-resource languages would probably have less of an incentive to utilize LLMs for editing as the results might not necessarily meet minimum standards depending on the level of representation of the language in the model. As such, the institutions that belong to particulate core universes associated with low-resource language might opt to prohibit the use of AI for writing and editing in their language. This could lead to these populations having a higher cost of academic text production than high-resource language populations and thus to an imbalance of power through knowledge production.

4.3 On Interlingual Knowledge Sharing

As expressed by Barros et al. 2023, there are some who hope that AI will serve as a bridge to unite the global south and the global north in matters of research and improve knowledge sharing. Particularly regarding translation and publishing where currently there are issues with inadequate attribution of authorship that arise from so-called honorary authors who are added despite providing minimal contributions that are often limited to language assistance (Barros et al. 2023). This has led some to make calls for the use of AI to fulfill this role and avoid this practice that is seen as detrimental to the process. Given this, I have selected interlingual knowledge sharing as the central issue for this section. In contrast to the first two sections of this chapter, interlingual knowledge sharing is not a topic that is seen by the literature presented as a potential issue but instead a current issue to which AI application is proposed as a solution. This type of AI implementation could be plausible in the sociotechnical imaginary constructed as it was stipulated that there would be a generalized conversational model that would have the capacity of code-switching and translating, both of which would be needed if intended to be applied in interlingual translation of specialized texts.

As per Barros et al. 2023, the process of interlingual knowledge sharing through publication could be described as so. A researcher or group of researchers develop an article in their native language but seek to publish it in a foreign journal. It is likely the case, given the examples provided by Barros et al. 2023, that these researcher know enough of the foreign language themselves to translate the article but are not confident in the absolute accuracy of their translation. So they seek a native speaker of the foreign language for assistance in correcting the paper to fit not only the foreign language requirements but also the nuances of the specific field of knowledge in that language community. To thank this foreign researcher the native researchers bestow upon them the title of honorary author and add them to the article they submit for publishing. Thus, Barros et al. say, that the practice dilutes the credit due to the real authors. If we were to model this process it would be like so. A researcher or group of researchers develop an article in their native language, that is, they produce a text while inhabiting a particular symbolic universe and several sub-universes of meaning from

where they draw the objectifications necessary to produce said article. But they wish to publish this article in a foreign journal, and thus they must ensure their article abides by the rules of a different language. As they are not native speakers of this foreign language, this language they have acquired through secondary socialization practices is not as cemented into their universe of meaning as their native language. They manage to translate it as they do at least partially inhabit the necessary universe of meaning but, since they do not fully inhabit it, they seek a native speaker of this foreign language for assistance. They ensure that this native speaker not only fully inhabits the necessary symbolic universe of the required foreign language but also any pertinent sub-universes of meaning, such as relevant fields of study or institutional affiliations, that would provide this researcher with the necessary vocabulary to fully assist them in the translation. This, they could not have done themselves and as such they bestow upon the foreign researcher recognition for enabling something that could not have been achieved by the researchers themselves. The foreign researcher on their part, as a native inhabitant of the language is able to grab the translation and polish it to better fit the requirements they are familiar with as per their own socialization practices. But the foreign researcher, not inhabiting fully the universe of meaning of the native language of the authors of the text, might not pick up on all the nuances of the original text and might make changes that, while fitting their own idiosyncrasy much better, might not necessarily correctly reflect the meaning of the original article.

A similar scenario but with the intervention of AI could go like so. A researcher or group of researchers develop an article in their native language and wish to publish it in a foreign journal. Since the proposed issue by Barros et al. 2023 mentions that the foreign researcher would be replaced by AI first we will follow the layout of the previous scenario. The researcher or group of researchers while only partially inhabiting the symbolic universe of the foreign language would then translate the text. Then they feed it into the AI and ask the AI to perform the same type of edits as the foreign researcher would have performed. Editing grammar mistakes, ensuring semantic accuracy, etc. In this manner, the processes made by the AI would almost mirror those described in section two of this chapter, that of editing processes. With the caveat of the fact that the researchers that would end up having to compare the output of the AI to their previously internalized expectations might not have the

knowledge or the confidence in their secondary socialized objectivations to detect with a critical eye whether the AI has done the task appropriately. This is thus a much riskier situation than simple editing of language mistakes because the end user might not be literate enough to judge the results independently. The result is then one of two options. The researchers might implicitly trust the output of the AI and present it as it is, incurring any risks of possible mistakes by themselves as AI has no accountability over the texts it produces. Or the researchers might then ask a foreign researcher to double-check the work of the AI, in which case it would not be much different from the original scenario proposed.

That was one option for the involvement of AI in translation processes of papers in academia. An alternative would be for the researchers to feed the original text in the native language to the AI and ask the AI to translate the text to the target foreign language. This scenario would be more similar to the writing issue explained in section two of this chapter than to editing with the constraint that the AI is not making up the words but “translating” them. In the sociotechnical imaginary constructed, it was stipulated that the generalized conversational model would have the capacity to code-switch and translate. However, the accuracy of the translations would still be up for questioning. Current examples of AI translation have shown that the models can make mistakes that in several instances have led to very serious real-world consequences (Bender et. al 2021: 618). The main issue identified by specialized researchers is that the overall coherence of the translated texts enable readers to miss cues that could have tipped them off to mistakes in the translation processes (Ibid). While it is possible that in the sociotechnical imaginary constructed these types of mistakes by the AI would be minimized, they would not likely disappear completely. So this scenario mixes elements of the high involvement scenario described in section two of this chapter, and the need for researchers to incur the risk of mistakes made by the AI in the previous scenario. An alternative would be to ask a foreign researcher to check the AI’s work, but with the added impediment that the words have not been chosen by the researchers themselves.

Overall, both of the options presented are similar in the end. Either the researchers risk making mistakes and assume the responsibility no matter whether they made those mistakes or if the AI did, or they ask for someone else’s help. There does not seem to be much of a difference in the processes with the involvement of AI with the possible exception of

increasing efficiency and with the caveat of having to deal with possible inaccuracies made by a machine incapable of taking accountability, which in itself is not a minor issue. This topic is very complex and as such there is no easy solution, but the application of AI in this instance seems to bring possible benefits that will entirely depend on the benchmarks that the models in the future achieve and whether they manage to convince researchers and the public at large that it is worth taking the risk because the machine is “accurate enough”.

But there is one topic that has not been mentioned yet. In the sociotechnical imaginary constructed, there is an asymmetry in the representation of languages based on the resources available. The previous examples were all contingent on the fact that both languages, source and target, would be high-resource languages. However, that is not the reality in all cases. If we consider the possibility of either of the languages being a low-resource language then we cannot say that involving AI would not make a difference. This is because while we might be able to speculate that the technology regarding AI translation would improve to the point where the risks are minimized, the sociotechnical imaginary mentions specifically asymmetries in language representation which lead to poor performance of the language model in low-resource languages. This would lead to three possible scenarios. One in which the low-resource language is the source language and the target language is a high-resource language. In this case, the mistranslation would lead to a poor representation of the work of the researchers and this would potentially perpetuate power imbalances. A second possible scenario is that in which the source language is the high-resource language and the target language is a low-resource language. In this case, the mistranslation could lead to the low-resource language-speaking community having access to inaccurate or poor-quality material which would again, lead to the perpetuation of power imbalances. Lastly, the third scenario would be a case in which both languages, source and target, are low-resource languages. In this case, both populations would be affected, one by poor representation and the other by access to low-quality material. Whichever the case if AI were to be employed for translation to or from a low-resource language the result would be quite different to that posited above as a probable net positive to society. Any low-resource language-speaking community would be affected by utilizing AI in this manner, and power imbalances would continue to be perpetuated.

4.4 On Interdisciplinary Research Efforts

In chapter one, it was mentioned that Ganguli et al. (2022: 1754) speak of the possibility of utilizing AI to finally address the constant friction present in interdisciplinary research. While not much is mentioned regarding actual implementation, it is possible to imagine a couple of alternative scenarios. Given that the sociotechnical imaginary constructed considers that there would be a generalized conversational model capable of code-switching, this last characteristic is the one that would be key to the plausibility of this particular endeavor. In order for AI to diminish the current friction present in interdisciplinary research, it needs to be able to switch between the codes of different disciplines, that is, to traverse between multiple sub-universes.

Interdisciplinary research has been gaining ever-increasing popularity within different academic communities due to its potential to address complex problems, despite it often presenting considerable difficulties in regard to the integration of diverse methodologies and theoretical frameworks. This is because the advantages and the disadvantages of interdisciplinary research both rely on the inherent differences present in the process itself. If we were to model interdisciplinary research as a whole, it would involve processes of externalization being conducted at the border of different sub-universes of meaning, the latter ones representing the different disciplines involved. Depending on how different the disciplines are and how many of them are involved, the process would get increasingly more complex. However, this description would only be applicable to a prototypical example, which is hardly akin to reality. In truth, interdisciplinary research would be better modeled from the perspective of individual researchers and not as a unified process. As such, the process would be like so: A researcher is participating in interdisciplinary research. From their point of view, their task is to perceive the objectivations produced by the other researchers, process them internally, and externalize their own objectivations subsequently. As the researcher internalizes the objectivations made by others, they need to be able to understand them, and for that, they need to at least partially reside within the others' sub-universe of

meaning. When externalizing their own objectivations, they need to, again, be aware of the others' sub-universe of meaning to attempt to produce objectifications that are understandable to the other researchers. It is evident that this process relies on a particular consciousness of the sub-universes of meaning inhabited by different participants by all the other participants. If every researcher participating is capable of performing this series of tasks repeatedly, then the result will be a success. But as is evident in the description, there are a plethora of pitfalls that are easy to fall into in the process, in particular, having multiple researchers inhabiting a series of different sub-universes would be difficult to execute and measure.

If we were to model interdisciplinary research with AI as a mediator, there are a number of alternatives for its implementation. However, the process would be somewhat simple, and different types of mediation could be modeled in a similar manner. That is, AI could be utilized to translate between the different sub-universes of meaning inhabited by different researchers. In the same way that it was described in section one of this chapter, but instead of searching its internal database or the Internet as it would in a literature search, the AI could provide an interpretation of the data provided by a particular researcher in the code of a different researcher. In this manner, the AI would be utilizing its inherent capacity to code-switch, stipulated in the sociotechnical imaginary, to cross the boundaries of the different sub-universes of meaning inhabited by the different participants to allow them to interpret the information in their own terms. While it might seem like this process could fall into a similar pitfall as the one described in the last section, that is, that the researchers would not have the capacity to adequately evaluate the quality of the output of the AI, this case is slightly different. Because code-switching is occurring between sub-universes of meaning that share a common core universe, that is, a common language-speaking community, it would be relatively easier for the researchers (that provide the data for the AI to interpret), to evaluate the quality of the output. And with this being the case, the difficulties apparent in the previous section are lessened significantly.

Of course, it is important to consider that this particular scenario modeled in this manner would only apply to interdisciplinary research being conducted among members of the same linguistic-speaking community. Only as long as all the participants are utilizing the AI in the manner described above, making sure to check that the interpretations made by the

AI have no significant deviation from the ones they would make themselves. These types of precautions could easily be codified by institutions to ensure that the appropriate care is being employed in the research process.

4.5 On the Epistemological Concerns of Outputs

One of the main concerns voiced by researchers in different texts presented in chapter one was that of the possibility of plagiarism originating from the use of AI. According to Livberber & Ayvaz 2023, this issue is perceived not only as originating from the possibility of researchers attributing texts produced by the AI to themselves but also the general understanding of the lack of a recognition of the place that LLMs have or can have in research environments. On the other hand, Bender et al. 2021 mention the general concern of the epistemological standing of LLMs outputs as these models are not capable of NLU but are often perceived as if they were. I have put both of these issues together as they both deal with the texts produced by the LLMs, what they are and what happens when you interact with them as a user. However, given the possible scope of this section and its incredible importance, it would be significant to mention that the analysis of this section will limit itself to the more salient aspects of the issue as relevant to the other sections of this chapter and will not represent an in-depth understanding of this topic.

According to Bender et al. (2021), Spärck Jones (2004) stated that there were certain unspoken epistemological and methodological commitments that anyone engaging with AI needed to partake in. “Either i) we commit ourselves to a noisy-channel interpretation of the task [...], ii) we abandon any goals of theoretical insight into tasks and treat LMs as “just some convenient technology” [p.7], or iii) we implicitly assume a certain statistical relationship — known to be invalid — between inputs, outputs and meanings.” (Bender et al. 2021: 615). As mentioned before, the literature presented in chapter one repeatedly made use of the word “tool” to refer to AI. This suggests that the current ground points the sociotechnical imaginary in the direction of option number two, to regard AI as “just some convenient technology,” or a useful tool. This is not to say that specialized researchers are not

questioning the epistemological standing of AI outputs, but it does highlight that the academic community at large is less concerned with the issue. These last ones, however, are concerned with plagiarism and the possibility of users claiming authorship of outputs of AIs as mentioned by Livberber & Ayvaz 2023. This creates an issue because in order to claim authorship of AI outputs one needs to recognize those outputs as epistemologically valid texts, meaningful in their own right. But as the quote above states, there is no statistical relationship between inputs, outputs, and meaning. So what is then this apparent capacity of authorship being perceived by non-specialists?

Foucault (1998) wrote that authorship has evolved through time, and what it means to be an author has changed. Recently, authorship has turned into a function, one that depends on “a complex operation that constructs a certain being of reason” (Ibid: 213). The key word in that quote being “constructs.” The author function is thus defined as a construct that depends, among many other things, on the institutional system and social perceptions of what or who the author of a particular text is (Ibid: 216). This understanding of authorship, while somewhat in opposition to the academic complex at large and its expectations of attribution of ownership of texts, does answer the previous query. Non-specialists are concerned with plagiarism of LLM’s outputs because they perceive the AI to be the author of said texts, with all the accompanying institutional and social systems and frameworks of understanding that enable this conclusion. So, while the epistemological standing of the outputs is still in question, and whether they hold meaning or not in their inception is up for debate, the apparent fluency and eloquence of the outputs seem to be enough for people to perceive the AI as having at least a minimum level of authorship over its own outputs.

One possible explanation for this perception can be found in Benveniste’s text from 1971. In it, they say, “It is through language that man constitutes himself as a subject” (Ibid: 224). They also explain that language itself is only possible through the subjectification of the emitter and that a dialectic reality is that which defines both the emitter and the receiver by their mutual relationship to one another (Ibid: 225). That is to say, when we provide a language-based input to the AI, we are subjectifying ourselves and the AI. When the AI “answers” our query with a language-based response, the machine is validating our reality; it is engaging with the processes of reality maintenance as described in chapter three. We do not

question whether the AI has intent or meaning; we assume it because it is engaging in the same mechanisms that we ourselves employ every day. As such, it is natural that we would attribute authorship to its outputs since in the framework of the social construction of reality, the AI is occupying the same role and engaging in the same reality-maintenance mechanisms that any other individual in that society. While the underlying workings that enable the AI to do this are not capable of NLU, the results it produces, to us, are indistinguishable from those of any other fellow member of our society. Denying a certain level of authorship to these outputs would leave us to then question the validity of the output, which in itself rests on the sheer facticity of its fluency. Thus, it is probable that in our sociotechnical imaginary, the vast majority of the public would not question the epistemological validity of AI outputs, as that would mean questioning the mechanisms of reality maintenance they employ every day. In exchange, through time, people would probably learn that akin to every human, LLMs make mistakes. They are subject to mistranslations, misinterpretations, biases, power imbalances, and more. Thus, as a society, we would slowly acquire the critical assessment tools necessary for evaluating the validity of outputs. In the background, specialists will continue the task of further understanding what these outputs truly amount to. But everyone else would remit themselves to thinking of AI as just another convenient tool or maybe even more.

All of this suggests that the conclusions of previous sections that involve limiting the use of AI through regulations to ensure that knowledge creation processes are not affected might end up coming from a different perspective. If we consider the repeated mentions of plagiarism concerns present in the literature, it stands to reason that measures will be taken at some point to ensure that people cannot morally attribute to themselves the authorship of AI outputs. This would in turn also affect some of the previous conclusions in the other sections. Particularly that of section two of this chapter and the high impact of some editing processes. At the very least, institutions would ensure that high-involvement editing processes would be regulated out of use in order to limit plagiarism concerns.

Discussion

Five concerns or problematic areas, from those identified in chapter one as expressed by researchers, have been modeled and analyzed in this thesis. These include: data fetching, scientific writing and editing, interlingual knowledge sharing, interdisciplinary research, and plagiarism and epistemological concerns. The results vary by area. Some have been identified as minimal or innocuous changes while others are of ethical concern. This distinction is made based on the level of impact of AI intervention, which is determined by comparing the three moments of the dialectic of the imagined future versus the current present for each problematic area.

When it comes to data fetching, two different levels were identified: data acquisition and data collection. The involvement of AI in both of these processes was found to be innocuous as AI simply serves to replace other tools that are currently used. So, if utilized in the same manner, the impact on knowledge construction is minimal at worst. For writing and editing, three levels were posited. On minimal intervention such as spell checking and proofreading, AI intervention was found to be of minimal impact on the dialectic. On the other hand, for high intervention processes, the dialectic is significantly affected, and ethical concerns arise. On medium intervention processes there is no precise definite conclusion as different types of interventions will yield different results. Thus, for scientific writing and editing, it can be said that to diminish the possibilities of ethical issues, institutions need to regulate the usage of AI on official documentation or publication material. Each institution will thus determine their own level of acceptable AI involvement by evaluating the ethical risks against the efficiency provided by the AI.

In the case of interlingual translation for knowledge sharing, the results are mixed. This area of concern is one identified as currently problematic for which AI is suggested as a solution. It was determined that AI would have minimal impact in the sense that the issue would not be solved through AI involvement, and while efficiency and autonomy might be increased, the user might also need to accept responsibility for any mistakes (either made by them or by the machine). This is because in this case, the AI is merely replacing the language

aid provided by foreign researchers while being incapable of taking accountability for the output it produces. In addition, language asymmetry representation is incredibly important in this particular topic, as if any low-resource language is involved, the results are less than satisfactory for knowledge sharing and might contribute to perpetuating power imbalances. On the other hand, interdisciplinary research was found to be a promising area for AI implementation. This is because in contrast with language translation, the users in this instance are able to assess the correctness of the outputs themselves with a much higher accuracy.

Finally, in the case of plagiarism and epistemological concerns, it was surmised that non-specialists will care little about the epistemological standing of outputs as AI, and particularly, LLMs are language based. The use of language in every step of user interaction with AI ensures that humans perceive a certain level of legitimacy and thus authorship from the AI's outputs as language is one of the main reality-maintenance mechanisms in society. So, while AI may lack natural language understanding, its outputs are indistinguishable from human-generated content, leading to societal acceptance. The case of plagiarism is thus relegated to institutional regulation, be it from institutes, universities, journals, or governments themselves.

Conclusion

This thesis has thus explored the possible impacts that AI might have on the social construction of reality through academia. It aimed to explore the possibilities surrounding the question: How could large language models impact the social construction of knowledge in academia? Emphasis on the chosen word *could* and not *would*. As it never aimed to provide a definitive answer per se, but to explore the possibilities that stem from such a future. For this purpose, the current reality and the expected possibilities of AI held by experts and non-experts alike were laid out in chapter one. This chapter served as the ground of this thesis, and several problematic areas were identified; be it problematic for the implementation of AI in the future or currently problematic and suggesting AI as a solution. Then, in chapter two, the technique of speculative semiotics was explained. Starting with a short exploration of speculative research in general, then a more narrow view of speculative semiotics, and finally the speculation process that in itself constitutes the construction of the sociotechnical imaginary. In chapter three, the model used was laid out, in particular the three moments of the dialectic were explained as per the literature, that is: externalization, internalization, and objectivation. Finally, in chapter four, the issues or problematic areas identified in chapter one were analyzed by comparing the possible future constructed with the present, utilizing the model laid out in chapter three.

Several dimensions have been touched upon: literature review, data collection, writing, editing, translation, interdisciplinary research, and even the epistemological underpinning of AI outputs. Some of these are included as areas where possible issues might arise with AI implementation, whereas others are areas where there are currently issues and AI is seen as a possible solution to them. In some cases, the impact of AI was unnoticeable, at least from the point of view of the moments of the dialectic that make up the processes analyzed. I mentioned this because while the weight of the different moments of the dialectic involved might not seem to be affected, that does not mean that the process itself does not change. It might make it more convenient or accessible to some while creating a barrier for others who thrive under traditional methods. These cases were such as data collection, data acquisition,

low to mid-impact editing, and interdisciplinary research. In other cases, it was evident that the processes that make up the issues analyzed were affected. In these cases, a clear impact on the moments of the dialectic could be established, putting the standing of academic processes at risk. These include areas such as writing, high-impact editing, and translation. In the last section, the issue of the epistemological standing of AI outputs, the capacity of authorship of AI, and the issue of plagiarism in regard to researchers claiming ownership of AI outputs is addressed. It was surmised that while specialized researchers will continue to evaluate the issue of correspondence between inputs, outputs, and meaning, non-specialized researchers are less concerned and take the outputs at face value. The fluency and eloquence of LLMs' outputs and the language-based process with which we interact with AI, force the user to engage with the machine as they would with a fellow human. Making use of the same reality-maintenance mechanisms that they employ in everyday conversations and that serve to reinforce the social constructs that make up the fabric of our understanding of reality.

The type of analysis performed in this thesis, while allowing for the exploration of otherwise unreachable realities, seems to fall short on a couple of topics. In particular, it was unable to address the call of researchers to utilize AI to find biases in academic texts. This was a conscious omission when selecting topics for the analysis section, as there was not enough literature found that addressed this as a real possibility in the future avenues of development of AI, nor any papers that suggested that it was able to perform such a task. Thus, given the lack of ground upon which to speculate, the choice was made to not address this particular issue in this thesis. This issue, however, remains open for future exploration contingent on the existence of new research regarding the capacity of AI at identifying biases to be used as ground. In addition, because this thesis was written on the basis of a generalization in order of social magnitude, it lacks nuances at the individual level. It is mentioned that in the sociotechnical imaginary, the generalized conversational model would have a symmetrical representation of interests. And, as it was explained, imbalances of power are still represented in the AI of the sociotechnical imaginary as they will continue to be part of social reality. The sociotechnical imaginary mentions that there would be an ongoing struggle to keep biases in check, and thus a continuous development of this technology. So, while we might be able to say that there would be a symmetrical representation of interests as

a generalization at the social level for the purposes of analysis, it is an assertion that needs to be held with care at any smaller scale.

The manner in which this analysis was performed and the methodological tools used allowed for the construction of a sociotechnical imaginary to frame and inform the analysis. The consequences of this choice are, on the one hand, that this analysis was unrestricted and was thus able to explore the possible impacts and changes of AI implementation without being restricted by the limitations of current realities. This freedom enabled the analysis to position itself at the extreme of possibilities, painting a clear picture of what a particular future, the one created based on the present as a ground, would look like. But in order to perform its objective, this thesis needs to affect the becoming of the present, as per the speculative technique applied. For that, the reader is called to action; to question their own interactions with AI and where it might lead them. It is only by looping the speculation back to the present that any real change might be effected. As such, the avenues of future research are abundant as this thesis remains an exploratory attempt into an unknown reality. In particular, section five of chapter four speaks about the epistemological understanding of AI outputs and leaves the doors open to linguistic analyses of AI representations that showcase the social role that users assign to them regardless of their own perceptions. In addition, section two of chapter four mentions that there are a myriad of possibilities for the implementation of AI into editorial processes and while some are seemingly innocuous others are quite impactful; as such, it invites an inquiry into where is the line that marks the point where editorial changes made by AI start to impact knowledge construction to a significant level.

References

- AlZu'bi, Shadi; Mughaid, Ala; Quiam, Fatima; Hendawi, Samar 2023. Exploring the Capabilities and Limitations of ChatGPT and Alternative Big Language Models. *Artificial Intelligence and Applications* 2(1): 28-37. <https://doi.org/10.47852/bonviewAIA3202820>.
- Bakhtin, Mikhail 1981. *The Dialogic Imagination: Four Essays*. Texas: *University of Texas Press*.
- Barros, Amos; Prasad, Ajnesh; Śliwa, Martyna 2023. Generative Artificial Intelligence and Academia: Implication for Research, Teaching and Service. *Management Learning* 54(5): 597-604. <https://doi.org/10.1177/13505076231201445>.
- Bender, Emily M.; Gebru, Timnit; McMillan-Major, Angelina; Shmitchell, Shmargaret 2021. On the Dangers of Stochastic Parrots: Can Language Models be Too Big? . *FAccT '21: Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*. New York: Association for Computer Machinery, 610-623.
- Benveniste, Émile 1971. *Problems in General Linguistics*. Miami: *University of Miami Press*.
- Berger, Peter L.; Luckmann, Thomas 1966. *The Social Construction of Reality*. New York: *Penguin Books*.
- Foucault, Michel 1998. What is an Author? In: Faubion, James D. (ed.), *Aesthetics, Method, and Epistemology*. New York: The New Press, 205-222.
- Ganguli, Deep; Hernandez, Danny; Lovitt, Liane; DasSarma, Nova; Henighan, Tom; Jones, Andy; Joseph, Nicholas; Kernion, Jackson; Mann, Ben; Askell, Amanda et al. 2022. Predictability and Surprise in Large Generative Models. *FAccT '22: Proceedings of the 2022 ACM Conference on Fairness, Accountability, and Transparency*. New York: Association for Computer Machinery, 1747-1764.
- Google 2025. SEO Fundamentals. *Google Search Central*. <https://developers.google.com/search/docs>
- Halewood, Michael 2017. Situated Speculation as a Constraint on Thought. In: Wilkie, Alex; Savransky, Martin; Rosengarten, Marta (eds.), *Speculative Research: The Lure of Possible Futures*. London, New York: Routledge, 52-63. ISBN: 978-1-138-68836-0.
- Jasanoff, Sheila; Kim, Sang-Hyun (eds.) 2015. *Dreamscapes of Modernity: Sociotechnical Imaginaries and the Fabrication of Power*. Chicago: *University of Chicago Press*.
- Kant, Rishi 2025. OpenAI's weekly active users surpass 400 million. *Reuters* 20.02.2025. <https://www.reuters.com/technology/artificial-intelligence/openais-weekly-active-users-surpass-400-million-2025-02-20/>.
- Khan, Sikandar 2023. AI at Doorstep: ChatGPT and Academia. *Journal of the College of Physicians and Surgeons Pakistan* 33(10): 1085-1086. <https://doi.org/10.29271/jcsp.2023.10.1085>.
- Koubaa, Anis; Boulila, Wadii; Ghouti, Lahouari; Alzahem, Ayyub; Latif, Shahid 2023. Exploring ChatGPT Capabilities and Limitations: A Survey. *IEEE Access* 11: 118698-118721. <https://doi.org/10.1109/ACCESS.2023.3326474>.

- Livberber, Tuba; Ayvaz, Süheyla 2023. The Impact of Artificial Intelligence in Academia: Views of Turkish Academics on ChatGPT. *Heliyon* 9(9): e19688. <https://doi.org/10.1016/j.heliyon.2023.e19688>.
- Pinzoliths, Robert 2023. AI in academia: An overview of selected tools and their areas of application. *MAP Education and Humanities* 4: 37-50. <https://doi.org/10.53880/2744-2373.2023.4.37>.
- Savransky, Martin 2017. The Wager of an Unfinished Present: Notes on Speculative Pragmatism. In: Wilkie, Alex; Savransky, Martin; Rosengarten, Marta (eds.), *Speculative Research: The Lure of Possible Futures*. London, New York: Routledge, 25-38. ISBN: 978-1-138-68836-0.
- Schaul, Kevin 2024. Is this AI? See if you can spot the technology in your everyday life. *The Washington Post* 02.04.2024. <https://www.washingtonpost.com/technology/interactive/2024/what-is-ai/>.
- Spärck Jones, Karen 2004. Language Modelling's Generative Model: Is it Rational? Technical Report: *Computer Laboratory, University of Cambridge*.
- Thibault, Mattia 2022. Speculative Semiotics. *Linguistic Frontiers* 5(3): 1-9. <https://doi.org/10.2478/lf-2022-0012>.
- Voloshinov, Valentin Nikolaevich 1973. *Marxism and the Philosophy of Language*. New York, London: *Seminar Press*.
- Webb, Taylor; Holyoak, Keith J.; Lu, Hongjing 2023. Emergent Analogical Reasoning in Large Language Models. *Nature Human Behaviour* 7: 1526-1541. <https://doi.org/10.1038/s41562-023-01659-w>.
- Wei, Jason; Tay, Yi; Bommasani, Rishi; Raffel, Colin; Zoph, Barret; Borgeaud, Sebastian; Yogatama, Dani; Bosma, Maarten; Zhou, Denny; Metzler, Donald; et al. 2022. Emergent Abilities of Large Language Models. *Transactions on Machine Learning Research*. ISSN 2835-8856.
- Wilkie, Alex; Rosengarten, Marsha; Savransky, Martin 2017. Speculative Techniques. In: Wilkie, Alex; Savransky, Martin; Rosengarten, Marta (eds.), *Speculative Research: The Lure of Possible Futures*. London, New York: Routledge, 113-116. ISBN: 978-1-138-68836-0.
- Zewe, Adam 2023. Explained: Generative AI. *MIT News* 09.11.2023. <https://news.mit.edu/2023/explained-generative-ai-1109>.

Kokkuvõte

Tulevike modelleerimine: Keelemudelitest ja akadeemilisest teadmusloomest

Selle magistritöö eesmärk on uurida, kuidas võiksid suured keelemudelid (SKM, ingl *large language models, LLM*) mõjutada teadmusloomet akadeemilises kontekstis. Selle saavutamiseks konstrueerib töö sotsiotehnilise kujutluse tulevikust, milles akadeemiline töö on läbi põimunud AIga. Kujutus põhistatakse praegusesse diskursusesse, mis käsitleb SKMe ja TIid akadeemilises kontekstis. Kujutluse ülesehitamiseks kasutab töö spekulatiivset semiootikat ning analüüsib seda siis Bergeri ja Luckmanni kolme dialektikamomendi abil.

Töös tuvastati järgmised probleemsed valdkonnad, millele spekulatsioon põhistatakse: andmete hankimine (nt kirjanduse otsimine ja andmekogumine), teaduslik kirjutamine ja toimetamine, interlingvistiline tõlkimine ning interdistsiplinaarne uurimistöö. Lisaks käsitleb töö muresid, mis seostuvad plagiadiga ning SKMide väljundi epistemoloogilise staatusega. Need tahud võetakse tükkideks lahti ja modelleeritakse lähtuvalt kolmest dialektikamomendist: internaliseerimine, eksternaliseerimine ja objektiviseerimine. Võrreldes võimalikku tulevikku olevikuga, püüab töö mõista TI potentsiaalset mõju spetsiifilistele teadustööga seotud akadeemilistele protsessidele.

Töö lähtub Mattia Thibault' spekulatiivse semiootika käsitlest kui tulevikusuunalisest kommunikatsioonist. Seda täiendavad Michael Halewoodi arusaam paigapõhisest spekulatsioonist (*situated speculation*) ning Martin Savransky vaade pragmaatilisele spekulatiivsele uurimusele. Sellest lähtuvalt joonistatakse välja neljaastmeline protsess, mille eesmärk on sisse seada vestlus olevikuga tuleviku kaudu, mis oleks põhistatud aktuaalsesse reaalsusse. Eesmärk ei ole seletada tulevikku, vaid mõjutada selle kujunemist ehk saamist. Keelel on siinses töös võtmeroll, eriti kuna see on esmatähtis ka Bergeri ja Luckmanni mudelis. Keel võimaldab kolme dialektikamomenti, kätudes nii kogemuse varaaida, subjektiveerimiste kristalliseerimise meetodi kui ka objektivsuse sunnijõuna. Keel on ka üks peamisi tegelikkust hooldavaks ja säilitavaks mehhanismiks ühiskonnas. Seetõttu on keelepõhistel interaktsioonidel (nt TIga „tšättimine“) kaugeleulatuvad implikatsioonid

Non-exclusive licence to reproduce the thesis and make the thesis public

I, Andrea Barone Renolfi,

1. grant the University of Tartu a free permit (non-exclusive licence) to

reproduce, for the purpose of preservation, including for adding to the digital archives of the University of Tartu until the expiry of the term of copyright, my thesis

Modeling Futures: On Language Models and the Social Construction of Knowledge in Academia,

supervised by Ott Puumeister PhD;

2. grant the University of Tartu a permit to make the thesis specified in point 1 available to the public via the web environment of the University of Tartu, including via the digital archives, under the Creative Commons licence CC BY NC ND 4.0, which allows, by giving appropriate credit to the author, to reproduce, distribute the work and communicate it to the public, and prohibits the creation of derivative works and any commercial use of the work until the expiry of the term of copyright;

3. am aware of the fact that the author retains the rights specified in points 1 and 2;

4. confirm that granting the non-exclusive licence does not infringe other persons' intellectual property rights or rights arising from the personal data protection legislation.

Andrea Barone Renolfi

23/04/2025