Is Cognitive Science at Aarhus University just Psychology?

An exploratory analysis of the syllabus of the Cognitive Science bachelor's degree at Aarhus University

Akaran Sivakumar, AU720104
202208344@post.au.dk
Social and Cultural Dynamics, Cognitive Science, Aarhus University
Jens Chr. Skous Vej 2, 8000 Aarhus, Denmark
03/06/2024

Cognitive Science emerged in the cognitive revolution in the 1950s and 1960s as an interdisciplinary field of psychology, linguistics, artificial intelligence, anthropology, philosophy, and neuroscience. It has since been criticised multiple times for being heavily influenced by Psychology and Computer Science, thus losing its interdisciplinarity. To investigate this, the syllabus for the Cognitive Science bachelor's degree at Aarhus University was examined. The syllabus data comprised book chapter summaries, research article abstracts, and course catalogue descriptions. This data was subject to topic modelling, so we could identify the representation of the six disciplines. Latent Dirichlet Allocation (LDA) was employed to infer topics, which were then compared to the embeddings of the six disciplines using a BERT (Bidirectional Encoder Representations from Transformers) model. The findings indicated a strong similarity between the syllabus content and the encoding of psychology, suggesting an imbalance in disciplinary focus. However, due to the limited scope of this project, this would need to be investigated further, using lecture slides and classroom instruction, to form an educated approach ensuring that Cognitive Science at Aarhus University is living up to the beliefs of the cognitive revolution.

Keywords: Cognitive Science, Topic Modelling, LDA, BERT, cosine similarity, interdisciplinarity.

1. Cognitive Science

Cognitive Science emerged from the cognitive revolution of the 1950s and 1960s, which aimed to overcome the limitations of behaviourism by creating a unified interdisciplinary field. It is built on the belief that “thinking can best be understood in terms of representational structures in the mind and computational procedures that operate on the structures” (Posner, 1996). The field was envisioned as a collaboration between six main disciplines: psychology, linguistics, artificial intelligence, anthropology, philosophy, and neuroscience (see Figure 1) (Bechtel, 2008). Today, Cognitive science is defined as: “The study of cognition or intelligence, the study of the cognitive processes involved in acquisition and use of knowledge” (*Cognitive Science, n. Meanings, Etymology and More | Oxford English Dictionary*, n.d.).

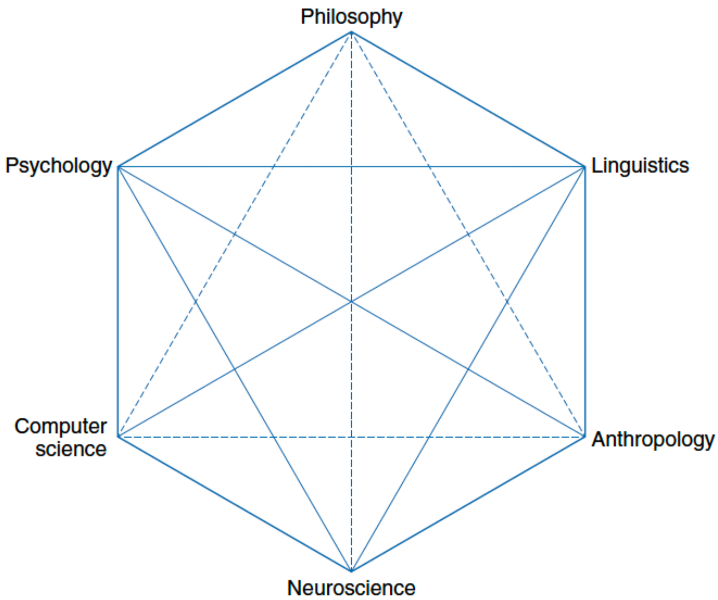

Figure 1 This figure shows the six disciplines that inspired the creation of the field of Cognitive science

In his book "Mind’s New Science," Howard Gardner noted that “dozens of scientists have attempted to define the nature and scope of the field” (Gardner, 1987). He argued that understanding human cognitive activities requires analysing both biological/neurological and sociological/cultural aspects. Margaret Boden suggests the field should be more accurately defined as: “*the interdisciplinary study of mind, informed by theoretical concepts drawn from computer science and control theory*” (Boden, 2006). The introduction of the journal Cognitive Science in 1977, which included "multidisciplinary" in its subtitle, highlighted this interdisciplinary aim. However, a 1998 report indicated that authors from psychology and computer science dominated the field, with less than 10% of affiliations from cognitive science departments (Von Eckardt, 2001). Today, the cognitive science faculties around North America are primarily training in the founding disciplines and some disciplines are much more present than others (Núñez et al., 2019). Consequently, Cognitive Science has evolved into a collection of diverse academic practices lacking common goals and paradigms.

1.1. Cognitive Science at Aarhus University

The cognitive science study programme at Aarhus University is unique, as it is the only bachelor's degree programme offered of its kind in Denmark. While Copenhagen University offers a similar programme, it is taught in Danish. The AU programme introduces students to fundamental theories of cognition and places focus on equipping them with the knowledge to design and carry out their investigations of the human mind, brain, and behaviour (*Bachelor's Degree Programme in Cognitive Science*, n.d.). This interdisciplinary programme exposes students to a broad range of concepts and methods, including experimental design, statistics, cognitive neuroscience, and cognitive approaches to communication and culture. Students learn about human decision-making, and how we use language to communicate, share feelings, and interact with others. Students will also learn computer programming (e.g. Python) and advanced tools for statistical data analysis (e.g. R and MatLab) which they will use in their experimental investigations of human cognition and behaviour.

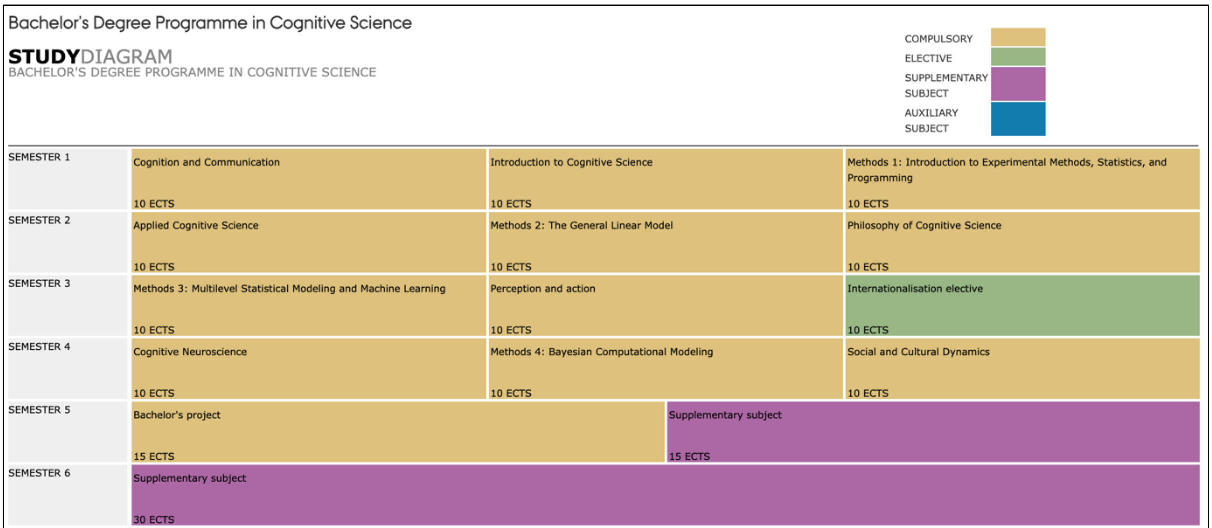

Figure 2 This figure contains the programme for BSc Cognitive Science at Aarhus University. The yellow courses are compulsory, the green is an elective course, and the purple courses are supplementary courses, that students can freely choose.

The study programme consists of 3 course types: Compulsory, Elective, and Supplementary subjects (*Bachelor's Degree Programme in Cognitive Science (2020)*, n.d.). Elective courses, an AU Arts initiative, are taught in English to foster an international environment (*Course Catalogue - Aarhus University*, n.d.). These electives can vary widely, as some students take them at other universities, during the summer, or even in another country. Due to this variability, the syllabi for elective courses have been excluded from the dataset. The same applies to the supplementary subjects in the 5th and 6th semesters, which also vary greatly and allow students to specialise in specific disciplines, making them hard to measure. This leaves the compulsory subjects, which all students must complete to earn a bachelor's degree in Cognitive Science. These courses cover various areas in the field, such as cognition and communication, philosophy of cognitive science, cognitive neuroscience, and methods courses that cover different aspects of statistical analysis and experimental design. In this investigation of the Cognitive Science study programme, we will only examine these compulsory subjects.

2. Topic modelling

Topic modelling algorithms are statistical methods that analyse the words of the texts in the corpus to discover the themes that run through them, how those themes are connected, and how they change over time. The magic of topic modelling algorithms lies in the fact that they do not require any prior annotations or labelling of the documents - the topics emerge from the analysis of the original texts. This allows for the organisation and summarisation of archives at a scale that would be impossible by brute force.

Latent Dirichlet Allocation (LDA) is one of the fundamental algorithms in topic modelling. The intuition behind LDA is that documents exhibit multiple topics, and it can capture these topics through statistics. It is most easily described by its generative process, an imaginary random process that the model assumes for the creation of the documents. We formally define a *topic* to be a distribution over a fixed vocabulary (set of words). For example, the *genetics* topic has words about genetics with high probability and the *evolutionary biology* topic has words about evolutionary biology with high probability. Each document reflects a mix of these topics, each with its distribution.

The goal of topic modelling is to uncover the topics from a collection of documents. While the documents themselves are observed, the topic structure - the topics, and the per-document per-word topic assignments - are *hidden/unknown (Blei & Jordan, 2003)*. This can be thought of as "reversing" the generative process to reveal the hidden structure. Topic modelling has proven to be useful in tasks such as information retrieval, classification, and corpus exploration.

2.1. Probabilistic topic models

Probabilistic topic modelling is a set of algorithms designed to find and label themes from massive datasets. LDA also falls under this category. In generative probabilistic modelling, we treat our data as arising from a generative process that includes *hidden variables*. This generative process defines a *joint probability distribution* over both the observed and hidden random variables. We can then use this joint probability distribution to calculate the conditional distribution of the hidden variables given the observed ones. This conditional distribution is the *posterior distribution (Blei et al., 2003; Steyvers & Griffiths, 2007).*

David Blei describes the posterior distribution mathematically as follows (Blei, 2012):

"We can describe LDA more formally using the following notation. The topics are denoted as $\beta_{1:K}$, where each β_k represents a distribution over the vocabulary. For the dth document, the topic proportions are represented by θ_d, where $\theta_{d,k}$ is the topic proportion for topic k in document d. The topic assignments for the dth document are z_d, where $z_{d,n}$ is the topic assignment for the nth word in document d. Finally, the observed words for document d are w_d, where $w_{d,n}$ is the nth word in document d, which is an element from the fixed vocabulary. With this notation, the generative process for LDA corresponds to the following joint distribution of the hidden and observed variables,

$$p(\beta_{1:K}, \theta_{1:D}, z_{1:D}, w_{1:D}) = \prod_{i=1}^{K} p(\beta_i) \prod_{d=1}^{D} p(\theta_d) \left(\prod_{n=1}^{N} p\left(z_{d,n} | \theta_d\right) p\left(w_{d,n} | \beta_{1:K}, z_{d,n}\right) \right)$$

Notice that this distribution specifies a number of dependencies. For example, the topic assignment $z_{d,n}$ depends on the per-document topic proportions θ_d. As another example, the observed word $w_{d,n}$ depends on the topic assignment $z_{d,n}$ and all of the topics $\beta_{1:K}$. (Operationally, that term is defined by looking up as to which topic $z_{d,n}$ refers to and looking up the probability of the word $w_{d,n}$ within that topic)."

By understanding what dependencies occur between topics, words, and topic assignments, LDA can categorise and interpret even the most complex data. For example, if we analysed a collection of news articles to identify themes, we could use the topic proportions and topic assignments, as well as the proportion of the themes within each article to find common themes such as politics, sports, and technology.

3. Methods:

3.1. Data Collection and Preprocessing

The data for this study was collected from the syllabus available on Brightspace on 20/05/2024. This syllabus included summaries of textbook chapters, research article abstracts and course catalogue descriptions from AU (*Bachelor's Degree Programme in Cognitive Science*, n.d.). The dataset comprised a total of 179 text entries. To prepare the data for analysis a preprocessing workflow was computed that involved several steps to remove noise, helping to create more accurate topic modelling.

1. **Text Cleaning:** First, extraneous whitespace, punctuation, and numerical digits were removed from the text. This was done to ensure that non-informative characters would not skew the results. This step is crucial for minimising the presence of non-informative characters that could skew the results (Krippendorff, 2019).
2. **Lowercasing:** All text was then converted to lowercase to maintain uniformity and ensure that no case-based discrepancies would occur (Bird et al., 2009).
3. **Stopword Removal:** Common stopwords, (e.g., "and", "the," "is"), were removed using the Natural Language Toolkit (nltk) library. This step was completed as these words do not carry any significant meaning. This will help focus the analysis on more meaningful words that contribute to the topic modelling process (Manning et al., 2008).
4. **Tokenization:** The cleaned text was then tokenised into individual words (tokens), thus preparing the text for vectorisation and analysis (Jurafsky & Martin, 2009).

All the mentioned steps were computed using Python's nltk library, as an attempt to follow the industry standard in computational linguistics (Loper & Bird, 2002). The implementation details, including parameter settings and code, are available in the GitHub repository attached to this project (Akaran19, 2024/2024).

3.2. Vectorisation and Topic Modelling with LDA

Following preprocessing, the text data was transformed into a document-term matrix using the CountVectorizer from the scikit-learn library (Pedregosa et al., n.d.). This process involves converting the corpus into a matrix of token counts, where each row represents a document (entry in the dataset), and each column represents a token from the vocabulary. The matrix entries correspond to the frequency of each token in each document (Harris, 1954). This representation is essential for enabling the application of Latent Dirichlet Allocation (LDA) to the text data.

Latent Dirichlet Allocation (LDA) is used to uncover latent topics within the corpus. This model allows for the discovery of seemingly abstract topics within collections of documents, making it one of the most powerful tools for thematic analysis of large text corpora (Steyvers & Griffiths, 2007). In this study, we specified the number of topics to be six, based on the vision behind the creation of Cognitive Science.

3.3. Keyword extraction and representative sentences

For each topic identified by the LDA model, the topmost representative 5 words were extracted. These keywords were selected based on their probabilistic associations with each topic, as determined by the LDA model. Extraction of these keywords serves as a summary of the themes captured by each topic (Chang et al., 2009). To contextualize each topic and enhance interpretability, representative sentences were identified from the original text corpus. These sentences were selected based on their high topic distribution scores, indicating their relevance to the corresponding topic.

3.4. Topic Naming with BERT

The challenge of naming topics accurately was addressed using BERT (Bidirectional Encoder Representations from Transformers). BERT is a transformer-based model designed to understand the context of words in a text (Devlin et al., 2019). The SentenceTransformer model from the sentence-transformers library was used to quantify the values in keywords and representative sentences (Reimers & Gurevych, 2019). This step allowed for BERT's contextual understanding to find the most representative words for each topic. The process involved the following steps: Embedding Generation, Centroid Calculation, and Similarity Measurement. To avoid duplication and ensure clarity, a check was implemented to ensure that the topic names were unique. Unique topic names aid in the ability to distinguish clearly between different themes. This involved maintaining a set of seen names and adjusting any duplicates. Unique topic names are crucial for effective communication of the results and to distinguish clearly between different themes (Blei, 2012). See Figure 3 for the whole workflow.

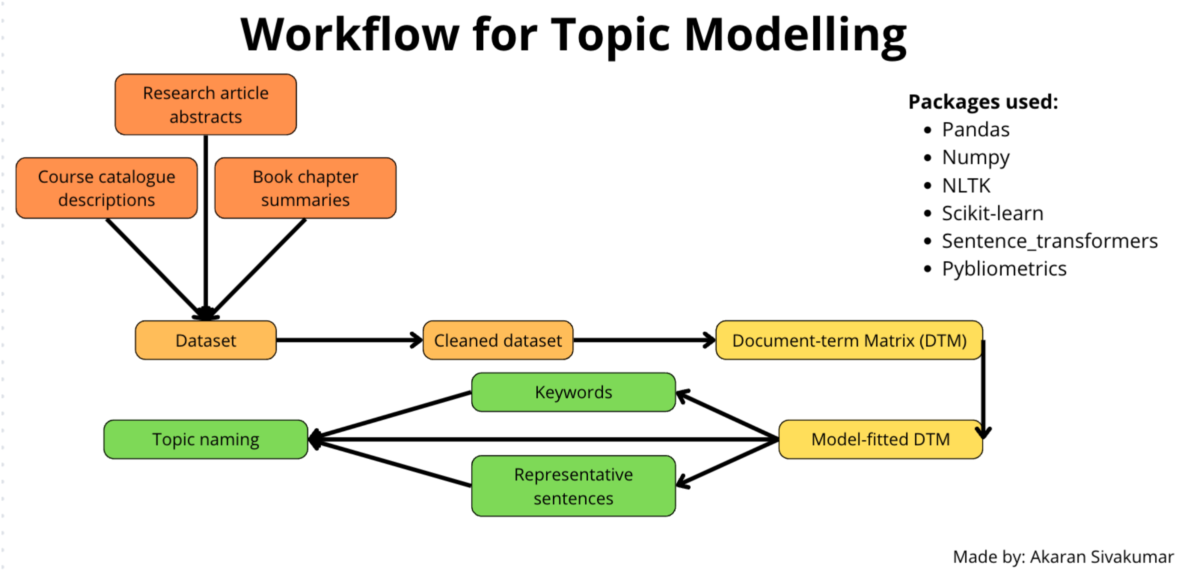

Figure 3 This figure presents the steps of topic modelling. It is, however, important to mention that the step where the topics are compared to the six disciplines is left out, as it was not a part of the same process.

3.5. Comparing topics to the disciplines of inspiration

The field of Cognitive science was inspired by the six disciplines of Psychology, Linguistics, Neuroscience, Anthropology, Artificial Intelligence/ Computer Science, and Philosophy. The original idea behind the cognitive revolution was to create this interdisciplinarity science that was comprised of all these six disciplines. To measure the interdisciplinarity of the study programme here at AU, we compare the similarity between the topics generated in the previous sections with the embeddings of the six disciplines in the model "all-MiniLM-L6-v2" from BERT (Devlin et al., 2019). The similarity was computed by using cosine similarity scores between topic embeddings and sentence embeddings generated by the model. We then compared these scores to an extended list of similar disciplines. The extended list included the following disciplines: Psychology, Neuroscience, Computer Science, Linguistics, Philosophy, Anthropology, Education, Mathematics, Statistics, Economics, Sociology, Biology, Medicine, Psychiatry, Neurology, and Communication Studies. The top 6 disciplines were then saved and 2 comparative radar charts were plotted using Seaborn and Matplotlib (Hunter, 2007; Waskom, 2021).

The similarity scores for the six original disciplines were also compared on a course-level, in an attempt to uncover where possible imbalances arise.

4. Results

Table 1 shows the topics generated by the topic modelling process. Each topic is described by the topic names encoded by BERT, the unique topic name chosen from the topic names, keywords extracted by the LDA model, and the most representative sentence chosen by their relevance to the topic.

Table of topics generated from Bsc Cognitive Science Syllabus				
Topic	Topic Names	Unique Topic Name	Keywords	Representative sentences
1	Research, data, inference	Research	Psychological, data, research, model	We provide a critical cumulative approach to vocal atypicalities in schizophrenia, where we conceptually and statistically build on previous studies.
2	Perception, action, visual	Perception	Stimuli, perception, effects, visual	Our ability to communicate relies in part upon a shared emotional experience, with stories often following distinct emotional trajectories and forming patterns that are meaningful to us.
3	Neurons, attention, sensory	Brain	neural activity, brain, sensory, information	Reading systematically activates the left lateral occipitotemporal sulcus, at a site known as the visual word form area (VWFA).
4	Memory, langauge, brain	memory	Memories, language, brain, human	Work with patient H.M., beginning in the 1950s, established key principles about the organization of memory that inspired decades of experimental work.
5	model, data, study	Modelling	Variables, linear, model, data	Here, we describe our motivations for developing experimental methods for studying cumulative cultural evolution and review the results we have obtained using these techniques.
6	Social, group, network	Networks	Coordination, network, social, groups	Recent work suggests that a population's ability to develop complex technologies is positively affected by its size and connectedness.

Table 1 This table shows the 6 topics that were created from the dataset. The topics are accompanied with prospective topic names, keywords, and representative sentences.

Figure 4 shows two subfigures. Figure 4a shows two different radar plots. The radar plots describe the similarity value of the dataset and the embeddings of six disciplines in BERT. The plot goes from 0 to 0.25 in all 6 dimensions, with each dimension being a discipline. The blue plot uses the six disciplines cognitive science was founded on, while the red plot uses an extended list of disciplines related to or connected with cognitive science. Figure 4b is a bar plot of the similarity scores calculated on a course level. This is done by extracting the text entries relevant to each course and then computing similarity to the six disciplines. Each of the courses is a different colour and the y-axis goes from 0 to 0.40, indicating that the similarity scores on a course level are higher than the whole syllabus.

a)

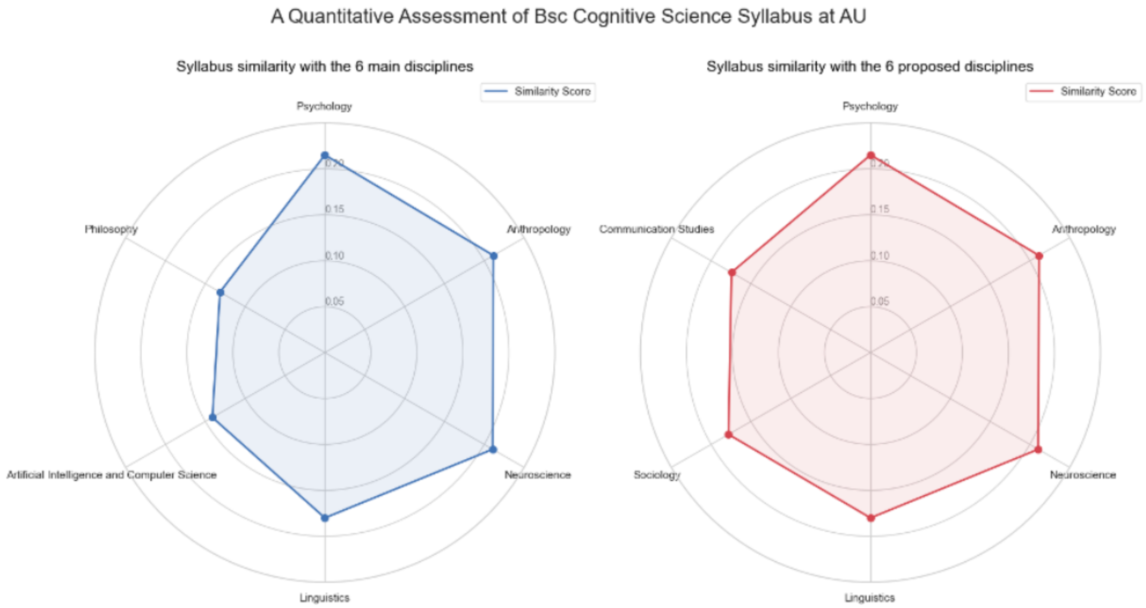

b)

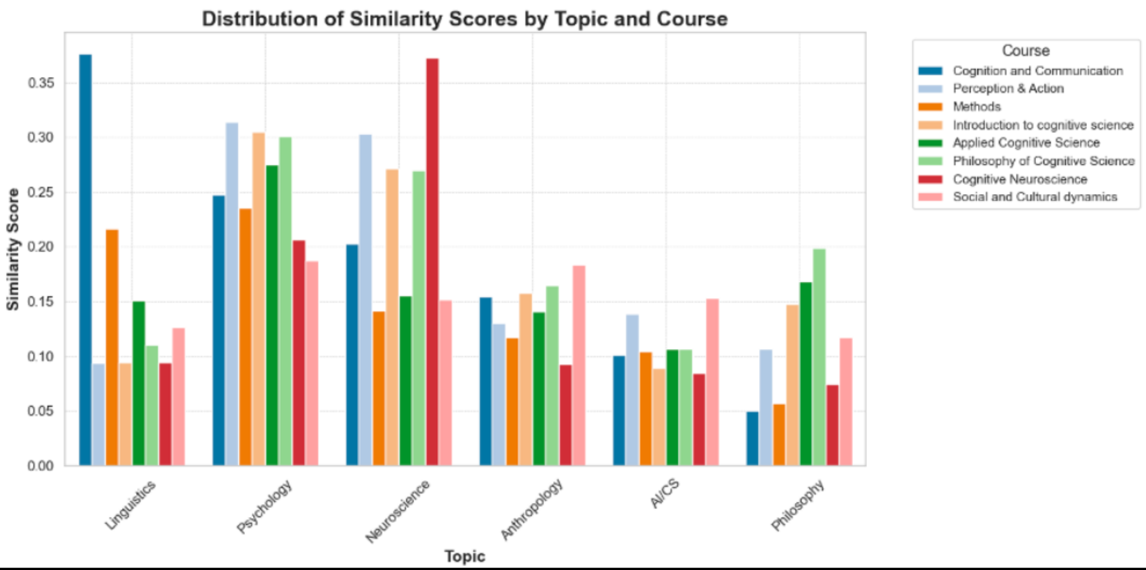

Figure 4 This figure contains the resulting plots from comparing the topics with the 6 disciplines. plot a) is composed of 2 subplots. the blue plot is the similarity for the syllabus with the six disciplines proposed in the cognitive revolution. The red plot is the similarity for the syllabus with 6 of the topics proposed in the list presented in the methods section.

5. Discussion

Table 1 shows the topics generated from the data. These topics are very similar, as many topics have the same keywords and exist in the same subdomain (psychology, statistics, etc.). Many of the topics also include “meta-words”, such as research, finding, effect, visual, psychological, and model. This could be due to a very relaxed preprocessing step, allowing many of these words to come through. But this could also be an artefact from the data that is used. The data comes from abstracts where researchers are forced to convey their results as concisely and broadly as possible. Otherwise, their papers will not appeal to the masses and thereby gain popularity. This leads to an exaggerated use of these “meta-words” as they convey much information for the space they take up. Possibly analysing the whole paper instead of just the abstract opens the opportunity for the intricate details of each paper to affect the topic modelling process, allowing for better specialization.

Figure 4a shows the similarity score for the data compared to the embeddings of six disciplines in BERT. The blue plot shows the six disciplines that cognitive science was supposed to be an interdisciplinary study of. Here, we can see that the dataset is not equally similar to all disciplines. Disciplines, such as Psychology, Anthropology, and Neuroscience score much higher than the rest, with especially Philosophy, Artificial Intelligence (AI) and Computer science (CS) scoring much lower than the rest. This signals some imbalance in the syllabus, that should be addressed if the study programme is to adhere to the principles of the cognitive revolution. This imbalance is further outlined in Figure 4b, where the similarity score is split up into a similarity score per compulsory course. Here we see that some courses succeed in providing a specialised deep dive in the specific discipline, such as the Cognitive Neuroscience course and the Cognition and Communication course. Other courses however fail to both provide specialized deep dives as well as providing an even amount of depth for each of the disciplines. The bar plot also further shows that Cognitive Science at AU is dominated by psychology, like the rest of the cognitive science study programmes mentioned in (Núñez et al., 2019). The second plot in Figure 4a is a revised radar plot with some other disciplines. This radar plot is more balanced, providing a guideline for what disciplines Cognitive Science at AU should market themselves as teaching. The radar plot suggests thinking of Cognitive Science as a balanced fusion of these six disciplines: Psychology, Anthropology, Neuroscience, Linguistics, Sociology, and Communication studies.

5.1. Future directions and limitations

This dataset is of course limited as we only had the written syllabus available. Also incorporating lecture slides or lecture recordings would provide a richer dataset more closely aligned to the study regulations and the reality of the study programme. We also do not know whether this similarity score is bad. To compare the score, we would have to compute similar scores for other study programs that have also emerged as an interdisciplinary field.

There were of course also limitations in the implementation of the methods. LDA is a very simple topic modelling approach. This leaves a lot of information "on the table" that could otherwise have been accessed and used to create better models. A very capable alternative to LDA is BERTopic (Grootendorst, 2022). BERTopic leverages BERT embedding to capture semantic relationships and contextual nuances among words, providing different and more accurate topic representations than the bag-of-words approach in LDA. The use of UMAP for dimensionality reduction also enhances the clustering process by preserving both local and global features of high-dimensional data (McInnes et al., 2020). BERTopic also leverages HDBSCAN's soft clustering approach, preventing unrelated documents from being misclassified (McInnes et al., 2017). Lastly, BERTopic's class-based TF-IDF method generates more meaningful topic-word distributions by emphasizing the importance of words within the clusters (Joachims, 1997).

6. Conclusion

In conclusion, the topic modelling of the syllabus at the bachelor's degree program of Cognitive Science at Aarhus University revealed an imbalance in how much the different disciplines influence the syllabus. The syllabus seems to be dominated by Psychology and Neuroscience, while disciplines such as AI/CS and Philosophy are barely present in the syllabus. Topic modelling with no prior topic names also showed a poor attempt at recreating the six disciplines, as "meta-words" were too prevalent, creating empty topics, such as research, modelling, and networks. The findings are however limited by the dataset, as this did not include any lecture slides or class exercises, both very meaningful in the training of cognitive scientists at AU. The method used was a LDA model with BERT aiding in the topic naming and similarity computing. LDA is a very simple model, with other methods such as BERTopic performing better in such circumstances. Future research should aim to address these limitations and explore new avenues of applying topic modelling in the research of study programmes. The findings suggest that the study programme does not live up to the idea of a uniform exploration of all six disciplines. To solve this, the syllabus should be revised to ensure all six disciplines are equally present and the study programme lives up to the ideals proposed by the cognitive revolution.

References

Akaran19. (2024). *Akaran19/soccult_exam* [Jupyter Notebook]. https://github.com/Akaran19/soccult_exam (Original work published 2024)

Bachelor's degree programme in Cognitive Science. (n.d.). Retrieved May 25, 2024, from https://bachelor.au.dk/en/cognitivescience

Bachelor's Degree Programme in Cognitive Science (2020). (n.d.). Retrieved June 2, 2024, from https://eddiprod.au.dk/EDDI/webservices/DokOrdningService.cfc?method=visGodkendtOrdning&dokOrdningId=16337&sprog=en

Bechtel, W. (Ed.). (2008). *A companion to cognitive science* (Nachdr.). Blackwell.

Bird, S., Klein, E., & Loper, E. (2009). *Natural language processing with Python*. O'Reilly.

Blei, D. M. (2012). Probabilistic topic models. *Communications of the ACM*, *55*(4), 77–84. https://doi.org/10.1145/2133806.2133826

Blei, D. M., & Jordan, M. I. (2003). Modeling annotated data. *Proceedings of the 26th Annual International ACM SIGIR Conference on Research and Development in Informaion Retrieval*, 127–134. https://doi.org/10.1145/860435.860460

Blei, D. M., Ng, A. Y., & Jordan, M. I. (2003). Latent dirichlet allocation. *J. Mach. Learn. Res.*, *3*(null), 993–1022.

Boden, M. A. (2006). *Mind as machine: A history of cognitive science*. Clarendon Press.

Chang, J., Gerrish, S., Wang, C., Boyd-graber, J., & Blei, D. (2009). Reading Tea Leaves: How Humans Interpret Topic Models. In Y. Bengio, D. Schuurmans, J. Lafferty, C. Williams, & A. Culotta (Eds.), *Advances in Neural Information Processing Systems* (Vol. 22). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2009/file/f92586a25bb3145facd64ab20fd554ff-Paper.pdf

cognitive science, n. Meanings, etymology and more | Oxford English Dictionary. (n.d.). Retrieved June 2, 2024, from https://www.oed.com/dictionary/cognitive-science_n

Course Catalogue—Aarhus University. (n.d.). Retrieved June 2, 2024, from https://kursuskatalog.au.dk/en?period=2&tags=iv-fag&year=2024

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2019). *BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding* (arXiv:1810.04805). arXiv. http://arxiv.org/abs/1810.04805

Gardner, H. (1987). *The mind's new science: A history of the cognitive revolution*. BasicBooks.

Grootendorst, M. (2022). *BERTopic: Neural topic modeling with a class-based TF-IDF procedure* (arXiv:2203.05794). arXiv. http://arxiv.org/abs/2203.05794

Harris, Z. S. (1954). Distributional Structure. *WORD*, *10*(2–3), 146–162. https://doi.org/10.1080/00437956.1954.11659520

Hunter, J. D. (2007). Matplotlib: A 2D graphics environment. *Computing in Science & Engineering*, *9*(3), 90–95. https://doi.org/10.1109/MCSE.2007.55

Joachims, T. (1997). *A Probabilistic Analysis of the Rocchio Algorithm with TFIDF for Text Categorization.* (p. 151).

Jurafsky, D., & Martin, J. H. (2009). *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition* (2. ed. [Nachdr.]). Prentice Hall.

Krippendorff, K. (2019). *Content analysis: An introduction to its methodology* (Fourth edition). SAGE.

Loper, E., & Bird, S. (2002). NLTK: The Natural Language Toolkit. *Proceedings of the ACL-02 Workshop on Effective Tools and Methodologies for Teaching Natural Language Processing and Computational Linguistics -*, *1*, 63–70. https://doi.org/10.3115/1118108.1118117

Manning, C. D., Raghavan, P., & Schütze, H. (2008). *Introduction to Information Retrieval*. Cambridge University Press; Cambridge Core. https://doi.org/10.1017/CBO9780511809071

McInnes, L., Healy, J., & Astels, S. (2017). hdbscan: Hierarchical density based clustering. *The Journal of Open Source Software*, *2*(11), 205. https://doi.org/10.21105/joss.00205

McInnes, L., Healy, J., & Melville, J. (2020). *UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction* (arXiv:1802.03426). arXiv. http://arxiv.org/abs/1802.03426

Núñez, R., Allen, M., Gao, R., Miller Rigoli, C., Relaford-Doyle, J., & Semenuks, A. (2019). What happened to cognitive science? *Nature Human Behaviour*, *3*(8), 782–791. https://doi.org/10.1038/s41562-019-0626-2

Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., & Cournapeau, D. (n.d.). Scikit-learn: Machine Learning in Python. *MACHINE LEARNING IN PYTHON*.

Posner, M. I. (Ed.). (1996). *Foundations of cognitive science* (5. print). MIT Pr.

Reimers, N., & Gurevych, I. (2019). *Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks* (Version 1). arXiv. https://doi.org/10.48550/ARXIV.1908.10084

Steyvers, M., & Griffiths, T. (2007). Probabilistic topic models. In *Handbook of latent semantic analysis.* (pp. 427–448). Lawrence Erlbaum Associates Publishers.

Von Eckardt, B. (2001). Multidisciplinarity and cognitive science. *Cognitive Science*, *25*, 453–470. https://doi.org/10.1016/S0364-0213(01)00043-X

Waskom, M. (2021). seaborn: Statistical data visualization. *Journal of Open Source Software*, *6*(60), 3021. https://doi.org/10.21105/joss.03021