1. Motivation

Foremost among the goals of K-12 education is that of equipping students with a foundational understanding of a wide array of knowledge domains. This foundational understanding grants students not only the freedom to later choose their areas of expertise, but also the ability to effectively communicate and collaborate with specialists in multiple fields -- a crucial concern in an increasingly interdisciplinary modern economy. However, the needs and opportunities of the modern economy are ever-growing, and as a result, so too are the number of potential areas of future expertise. This creates inevitable pressures for curriculum designers. Should students learn more chemistry in anticipation of increased opportunities in biotech, or should they learn more math and statistics in anticipation of the growing opportunities in Machine Learning? Ideally, we'd like them to learn more of both; but time is a limited resource, and the more is spent in one area, the less can be afforded to another.

One hypothetical means through this dilemma is to ease the burden of equipping students with a thorough understanding in any given field by determining the most efficient means to impart the central or foundational aspects of each field. Often, these central aspects are determined by committee -- experts in a given field form consensus as to what they believe is most important, provisional curriculums are designed to try to teach those topics in whatever order the experts suggest, educators report back on difficulties encountered, refinements are made, and the process is repeated until no further improvement seems viable.

How well this process ultimately works is a matter of debate. Evidence suggests that students forget forget academic content at a rate of 50% every two years[1], meaning most of what they've learned in K-12 is completely forgotten by the time they graduate. Worse still, where we would hope students take advantage of what they *do* retain to approach real world situations from a number of perspectives, we find that more often they often fail to use any cognitive tools which a problem doesn't explicitly call for. In short, students forget most of what they learn, and fail to **transfer** most of what they happen to retain.

We consider that part of the problem might be inherent to the curriculum design process. When experts in a given field determine that some topic is important, they do so from their experience *as experts in that field*. But, while we do hope that students eventually attain mastery in at least one field, we don't intend for them to become experts in *every* field we introduce them to. And what is of prime import for eventual mastery is not necessarily of any benefit for generalist use. To further complicate matters, experts are also working against the "Curse of Knowledge", in which their understanding makes them blind to the information non-experts may be missing. In these cases, Core Concepts may be so internalized by experts as to be deemed too trivial focus on, or overlooked altogether. Students are then left to slowly work backwards, triangulating Core Concepts through sufficient exposure to the otherwise uninteresting or inscrutable topics they inform.

We consider alternative approaches. It may be the case that by prioritizing only Core Concepts which are most relevant to *all other* concepts in a given field, we can satisfy the true foundational requirements of each field in the fewest number of concepts, thereby reducing cognitive and curricular load. It may also be the case that by identifying concepts which are most relative to multiple fields, we may equip students to be more agile generalist outside of their chosen domain of expertise, giving them the breadth they need to identify and dive deep into only the small portion of advanced knowledge they may need for a particular novel task. As an added benefit, by prioritizing concepts from each field which are most relevant to concepts outside of that field, we can design curricula where concepts are routinely used across domains -- in effect granting students more practice with thinking flexibly and approaching a problem from multiple perspectives

For any of these alternative approaches to be viable, some system is required by which Core Concepts and their interrelations can be identified from real world data, or by which their interdependencies can be honed in on within existing curricula. The potential benefits of a system approximating this functionality go beyond the motivations that inspire this research, and the potential uses of such a system are diverse enough that we feel periodic reappraisal of what the state-of-the-art in Machine Learning allows for is well worth the effort. Section 2 of this paper outlines our explorations into automated topic generation and categorization from loosely structured course materials across competing curricula. Potential uses, and an appraisal of current technical limitations is offered in section 3.

2. Methodology

A fully connected Neural Network with a single hidden layer[2] of 300 neurons was trained on introductory Physics course material comprising Crash Course lectures and The Feynman Lectures. The network was fed lectures in small, roughly sentence-length text excerpts. Each excerpt had one word omitted, and the network was trained to fill-in-the-blank by selecting from a list of 10,737 possible candidate words. Once training was complete, the resulting weights between each of the 10,737 input layer neurons (where each neuron corresponded to a unique word in the corpus) and the 300 hidden layer neurons were treated as coordinates embedding each word in the corpus in a 300-dimensional euclidean space. By the constraints of the training procedure, this embedding places related words closer to one another in the 300-dimensional space, and unrelated words further from one another in such a way that directions in the space roughly correspond to semantic relationships between words[2].

In order to limit the dataset to topical words without removing potentially informative semantic information, common words were determined and filtered out post-training by frequency analysis of the physics corpus against a more general corpus. This reduced the dataset down to 855 words and mathematical symbols. Clusters of related words were then identified through a fuzzy k-means clustering algorithm[3]. Clusterings were generated from k=2 up through k=9 using a cosine-similarity distance metric. The identified clusters were then manually reviewed and tagged to best describe the most representative words in each cluster at each clustering level.

Finally, three parameterization of t-SNE were run to reduce the 300-dimensional embedding to three 3-dimensional visualizations roughly preserving the relative cosine distances between data-points.

3. Results and Future Work

Qualitative appraisal of clustering results indicate clearly discernible subtopics through all explored clustering levels. However, the distances between the centroids of these topics sheds little light on their inter-relations. The embedding is properly descriptive, but lacks any clear insights of prescriptive value.

The descriptive possibilities are intriguing. By clustering the embedded words, the system is in some sense taking raw course material and categorizing it into topics or modules. This should in theory allow for automated reconciliation of multiple curricula without human intervention. In other words, sequences through sub-topics can be automatically inferred from one curriculum and compared against or extend other automatically inferred sub-topic sequences from competing curricula on the same subject. If a database of known sub-topic sequences is created by this process, a directed acyclic graph can be generated from the differences in these sequences, which -- assuming sufficient robustness -- could then be leveraged to multiple hypothetical ends. For example:

- *The unique educational history of a transfer student could be analyzed, and the exact sub-topics the student might require in order to effectively progress in a new curriculum could be automatically inferred.*
- *Average student performance on any one sub-topic can be compared between competing curricula, and determinations could be made as to the relative value of each curriculum for mastery of a given sub-topic. Over a sufficient number of curricula, determinations could in principle be made as to which sub-topics serve as the most critical bottlenecks in need of further attention.*
- *Proposed unstructured course material can be automatically inspected to determine potential prerequisite sub-topics for self-directed learners or teachers considering alternative or supplemental resources.*

- *Curriculum designers can select "goal" sub-topics, and an existing curriculum can be checked for redundancies or shortcomings with respect toward those goals, or new curricula can automatically be generated inferring the shortest path to those goals from a graph of known viable paths.*

The prescriptive prospects are, unfortunately, not as promising. The system in its current state does little to aid identification of concepts which may be truly "Core," or concepts which may offer the most benefit in terms of knowledge transfer. This is in some small part due to the limited size and diversity of the corpus, and in much larger part due to the limited potential for understanding of the chosen network architecture.

This last point is the most troublesome as even current state-of-the-art methods are unlikely to offer additional benefit for our purposes[4, 5]. Automated identification of high-centrality or high-transfer concepts may remain unfeasible until the state-of-the-art in Natural Language Processing matures beyond inference of semantic relations, and -- at the very least -- to more sophisticated inferences of *ontological* relations.

It is worth noting that there are some indications that too high a dimensional space for the embedding was chosen, and further hyperparameter tuning may be beneficial. Additionally, little attention was given to information encoded as principal directions in the embedding, and interested readers are welcome to carry out their own analysis of the data.

The original 300-dimensional embedding, as well as the corresponding words and 3d representations, may be accessed from the URLs provided in the appendix.

References

[1] Subirana, B., Bagiati, A., & Sarma, S. (2017). *On the Forgetting of College Academics: at "Ebbinghaus Speed" (68), 1–12.*

[2] Tomas Mikolov, Wen-tau Yih and Geoffrey Zweig. *Linguistic Regularities in Continuous Space Word Representations. In Proceedings of NAACL HLT, 2013.*

[3] C.J. Bezdek, *Pattern Recognition with Fuzzy Objective Function Algorithms, Plenum Press, New York, 1981*

[4] Zewei Chu, Hai Wang, Kevin Gimpel, and David McAllester. *Broad context language modeling as reading comprehension. In Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers, volume 2, pages 52–57, 2017.*

[5] Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. 2018. *Universal transformers. preprint arXiv:1807.03819*

Appendix

300-dimensional embedding: https://curriculumredesign.org/wp-content/themes/iloChildTheme/data/tensorCleaned.csv

Words and 3-dimensional representations: https://curriculumredesign.org/wp-content/themes/iloChildTheme/data/consolidated_projections.csv