

Toward a Thermodynamics of Meaning

Jonathan Scott Enderle^a

^aUniversity of Pennsylvania Libraries, 3420 Walnut St., Philadelphia, PA 19104-6206, United States of America

Abstract

As language models such as GPT-3 become increasingly successful at generating realistic text, questions about what purely text-based modeling can learn about the world have become more urgent. Is text purely syntactic, as skeptics argue? Or does it in fact contain some semantic information that a sufficiently sophisticated language model could use to learn about the world without any additional inputs? This paper describes a new model that suggests some qualified answers to those questions. By theorizing the relationship between text and the world it describes as an equilibrium relationship between a thermodynamic system and a much larger reservoir, this paper argues that even very simple language models do learn structural facts about the world, while also proposing relatively precise limits on the nature and extent of those facts. This perspective promises not only to answer questions about what language models actually learn, but also to explain the consistent and surprising success of cooccurrence prediction as a meaning-making strategy in AI.

Keywords

language modeling, natural language semantics, artificial intelligence, statistical mechanics

1. Introduction

Since the introduction of the Transformer architecture in 2017 [29], neural language models have developed increasingly realistic text-generation abilities, and have demonstrated impressive performance on many downstream NLP tasks. Assessed optimistically, these successes suggest that language models, as they learn to generate realistic text, also infer meaningful information about the world outside of language.

Yet there are reasons to remain skeptical. Because they are so sophisticated, these models can exploit subtle flaws in the design of language comprehension tasks that have been overlooked in the past. This may make it difficult to realistically assess these models' capacity for true language comprehension. Moreover, there is a long tradition of debate among linguists, philosophers, and cognitive scientists about whether it is even possible to infer semantics from purely syntactic evidence [26].

This paper proposes a simple language model that directly addresses these questions by viewing language as a system that interacts with another, much larger system: a semantic domain that the model knows almost nothing about. Given a few assumptions about how these two systems relate to one another, this model implies that some properties of the linguistic system must be shared with its semantic domain, and that our measurements of those properties are valid for both systems, even though we have access only to one. But this conclusion holds

CHR 2020: Workshop on Computational Humanities Research, November 18–20, 2020, Amsterdam, The Netherlands


✉ enderlej@upenn.edu (J.S. Enderle)

🌐 <https://senderle.github.io> (J.S. Enderle)

🆔 0000-0003-1901-7921 (J.S. Enderle)

© 2020 Copyright for this paper by its authors.

Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

only for some properties. The simplest version of this model closely resembles existing word embeddings based on low-rank matrix factorization methods, and performs competitively on a balanced analogy benchmark (BATS [9]).

The assumptions and the mathematical formulation of this model are drawn from the statistical mechanical theory of equilibrium states. By adopting a materialist view that treats interpretations as physical phenomena, rather than as abstract mental phenomena, this model shows more precisely what we can and cannot infer about meaning from text alone. Additionally, the mathematical structure of this model suggests a close relationship between cooccurrence prediction and meaning, if we understand meaning as a mapping between fragments of language and possible interpretations. There is reason to believe that this line of reasoning will apply to any model that operates by predicting cooccurrence, however sophisticated. Although the model described here is a pale shadow of a hundred-billion-parameter model like GPT-3 [5], the fundamental principle of its operation, this paper argues, is the same.

2. Previous Work

Most recent work on language modeling builds on the word2vec word embedding model [20] and its descendants such as GloVe [23]. These models drew from a longer tradition of distributional semantics in linguistics [11] [7] and early machine translation research [30] [18] [17] [8]. The promise of word embedding models for research in the humanities was quickly recognized, leading to historical studies of analogical language [12] and diachronic lexical change [10], but questions remained about the utility of embeddings for close humanistic analysis. Word embedding models suffer from stability problems, yielding seemingly precise answers that change when training input is modified only slightly [1], and their internal geometric structure is poorly understood [21].

Attempts to build a better theoretical understanding of word embeddings have often focused on exploring the ways different models prove to be mathematically equivalent in some limit [15] [2], or showing the importance of preprocessing and hyperparameter selection. In many cases, with optimal hyperparameter choices, factorizing word cooccurrence matrices using SVD and a log weighting is sufficient to produce results competitive with state-of-the-art models [16] [9]. For these reasons, the claim that word embeddings are indeed representations of meaning, and not merely dense representations of word cooccurrence, still lacks strong theoretical support. On the other hand, even simple cooccurrence data seems intuitively to capture something about meaning in a way that remains mysterious [2].

More recent language modeling has focused on sequence prediction, either using recurrent neural networks [24] or attention-based mechanisms [29]. Large language models using the Transformer architecture apparently capture rich semantic information usable in a range of downstream applications [13]. But as with word embeddings, there remain empirical and theoretical reasons to be skeptical that these models are capturing information about meaning, rather than performing an extremely sophisticated and accurate version of positionally-aware cooccurrence prediction. At least some attempts to use Transformer models to perform challenging natural language comprehension tasks have shown that existing problem datasets contain subtle linguistic cues that leak information about correct answers [22] [19]. These cues have been missed in the past, but with their linguistic sophistication, newer models recognize them, producing spurious state-of-the-art results without demonstrating true comprehension.

Recent work by Bender and Koller [4] provides an even stronger theoretical case against the

claim that language models infer meaning beyond simple cooccurrence. Synthesizing arguments and evidence from linguistics and philosophy, including Searle’s famous Chinese Room argument [26], Bender and Koller argue that “the language modeling task, because it only uses form as training data, cannot in principle lead to learning of meaning.” Or, in Searle’s pithy formulation, the operations of a computer have “syntax but no semantics.” Bender and Koller’s reliance on Searle is notable, given that Searle’s argument was not against language modeling, but against the very possibility of artificial intelligence. Anyone who takes his reasoning entirely seriously should forever abandon the notion that a computational process could truly comprehend meaning. Yet in their final analysis, Bender and Koller back away from Searle’s strongest claims, acknowledging that “if form is augmented with grounding data of some kind, then meaning can conceivably be learned to the extent that the communicative intent is represented in that data,” and that a sufficiently successful language model “has probably learned something about meaning.”

3. Meaning, Cooccurrence, and Thermodynamics

How can we synthesize these seemingly contradictory bodies of theory and evidence? It’s plausible to claim that language models can never do more than predict the way elements of language cooccur in text, since they never see any other kind of evidence. And yet even the simplest kinds of cooccurrence prediction, such as basic matrix factorization, produce surprisingly good representations of something that looks intuitively like meaning. Suppose that rather than examining the details of particular language models to see how they differ, and which might be more or less correct, we focus on what they have in common. Is there some unrecognized connection between meaning and cooccurrence prediction in all its forms?

This section proposes such a connection based on a model borrowed from statistical mechanics. Similar approaches have been applied to practical language modeling problems [28] [27] and theoretical discussions of algorithmic and semantic information [3] [14]. But to the author’s knowledge, no prior work has used thermodynamic analogies to specifically investigate the relationship between language and its semantic domain.

This model begins by treating interpretations as possible configurations of an unknown physical system. It then constructs a statistical mechanical partition function that counts the number of interpretations applicable to each fragment of language in a corpus. It immediately follows that the Hessian of that function—the matrix of its mixed second partial derivatives—is a covariance matrix describing word cooccurrences. The Hessian matrix can be used, in turn, to approximate directional derivatives of the partition function, which describe the ways the partition function changes when the meanings of words are slightly modified. These directional derivatives are word vectors, with all the expected properties.

3.1. Model Assumptions

Setting up our model requires some odd assumptions about how language works. To begin with, it requires that we assume that meaning is quantifiable in the most naive way. It’s not uncommon in colloquial speech to talk about the amount of meaning a phrase has, without specifying what the phrase means. Some phrases, we might say, are meaningless; others are full of meaning. To construct a statistical mechanical model of meaning, it is useful to assume that this is a perfectly correct way of quantifying meaning, and that, so quantified, meaning is a conserved value that plays the same role as energy in a typical thermodynamic ensemble.

As long as we are making extravagant assumptions, let's also assume that for a given linguistic system and an associated semantic domain, words have a stable average capacity for holding meaning, and that word counts are conserved values just like energy, so that a combined linguistic system and associated semantic domain contains an unknown but fixed number of copies of every possible word. Leaving aside the linguistic significance of these assumptions for a moment, we can skip ahead by recognizing them as formally equivalent to the assumptions made in the construction of the grand canonical ensemble.

3.2. The Grand Canonical Ensemble

In classical thermodynamics, the grand canonical ensemble describes a system of particles—such as a container of gas—that is in thermodynamic and chemical equilibrium with a much larger system, a reservoir of energy and particles. Concretely, this means that the temperature of the gas in the container is the same as its surroundings (assumed to be homogenous, and far larger than the container), and that the container can exchange particles with its surroundings, but at a steady state, so that it is as likely to lose a particle as to gain a particle at any given moment. Furthermore, both the amount of energy and the number of particles shared between the container and its surroundings are fixed—energy and particle number are conserved values.

To understand the behavior of this ensemble, we begin by imagining that we could track the exact position and momentum of every particle in the system (container), as well as the exact position and momentum of every particle in the reservoir (surroundings of the container). At a given instant in time, these values constitute a “microstate.” Since we have both a system and a reservoir, we can divide a single microstate into parts, considering just the microstate of the system, or just the microstate of the reservoir. We can also determine that certain system microstates are incompatible with certain reservoir microstates, because the combination would violate a conservation law. In other words, for some pairs of system microstate and reservoir microstate to coexist, energy or particles would have to be created or destroyed, which would violate the rule that energy and particle number are conserved values.

If we rule out all system-reservoir microstate pairs that are not compatible (\nleftrightarrow), and assume that all reservoir microstates are equally likely—an acceptable approximation when the reservoir is far larger than the system—then we can approximate the probability of a given system microstate s_i by counting the number of reservoir microstates that are compatible (\leftrightarrow) with it. Using Iverson brackets ($[i = j] = \delta_{ij}$) we can say

$$p_i \propto \sum_j [r_j \leftrightarrow s_i] \tag{1}$$

To recover the probability itself, we can divide by the sum over all s :

$$p_i = \frac{\sum_j [r_j \leftrightarrow s_i]}{\sum_{j,k} [r_j \leftrightarrow s_k]} \tag{2}$$

These sums are very large, and it's not clear how to calculate them. But it turns out we don't need to. Given a few standard assumptions from thermodynamics, our assumptions about conserved quantities, and a bit of calculus, it's possible to use them to derive the following function:

$$Z = \sum_i e^{\beta(\mu N_i - E_i)} \quad (3)$$

This is the grand canonical partition function. From it, we can then directly calculate the probability of system microstate i like so:

$$p_i = \frac{e^{\beta(\mu N_i - E_i)}}{Z} \quad (4)$$

This formula tells us, first, that at a given fixed temperature T determined by $\beta = \frac{1}{k_B T}$, system microstates with more energy (E_i) are less probable, because they are compatible with fewer reservoir microstates. It also tells us that for any given energy level, system microstates containing more particles (N_i) with a higher chemical potential (μ) are more probable. This is because given two systems with the same energy, the system with a higher overall potential has a higher energy capacity.

This partition function can be extended to systems that have multiple kinds (“species”) of particles. In that case, each species has its own chemical potential and count. For a system with k different species

$$Z = \sum_i e^{\beta(\mu_1 N_{1,i} + \mu_2 N_{2,i} + \dots + \mu_k N_{k,i} - E_i)} \quad (5)$$

$$p_i = \frac{e^{\beta(\mu_1 N_{1,i} + \mu_2 N_{2,i} + \dots + \mu_k N_{k,i} - E_i)}}{Z} \quad (6)$$

3.3. From Compatibility to Interpretation

What does all this have to do with language? The first hint that the grand canonical partition function might have some usefulness as a model for language is that energy and meaning (in the naive quantitative sense described above) both impose similar compatibility constraints on the system and reservoir. Just as higher-energy states in the system correspond to fewer possible reservoir states, more meaningful sentences correspond to fewer possible interpretations. A statement with less meaning has less precision, while a statement with more meaning has more precision, eliminating a larger number of possible interpretations. The line of reasoning is similar for particle species and words. Just as a particle species with higher chemical potential has higher energy capacity, a word with a higher “semantic potential” has a higher capacity for meaning.

Consider, for example, the sentence “It stinks.” Then compare it to “On January 15, 2008, a rainfall of 110mm was recorded in the city of Dubai.” The specific interpretations these sentences can be given will depend on context, and in some contexts, “It stinks” might be a meaningful and precise sentence. But on balance, we should expect “It stinks” to be less meaningful than “On January 15...,” both because it contains fewer words, and because the words it contains are less precise than words like “rainfall” and “Dubai.”

Although this is a simple way of thinking about meaning, it is not as simplistic as it may seem at first. Consider the sentence “Ask for me tomorrow, and you shall find me a grave man,” as uttered by a dying Mercutio. One might think that by the logic above, this sentence would be made less meaningful by the presence of an ambiguous word, “grave,” here meaning either “serious” or “a place of burial.” But a more careful analysis leads to a different conclusion. If these two senses were available independently, and the sentence could be properly interpreted

in two different ways, it would indeed be less meaningful because of this ambiguity. In this context, however, choosing just one of those senses to the exclusion of the other would yield a misreading of the sentence. It does not invite two different possible interpretations; it invites one interpretation that combines together two distinct concepts both conveyed by the word “grave.” By eliminating interpretations that do not combine these two senses together, this sentence uses ambiguity to achieve a higher degree of precision. Analyzed this way, literary language is often likely to be more precise and meaningful than everyday language, despite sometimes having greater surface ambiguity.

If we translate these ideas into a mathematical form, and start thinking about compatibility (\leftrightarrow) as a semantic relationship, then equation 2 says roughly that the probability of a given sentence (s_i) is equal to the number of interpretations ($r_{j,\dots,k}$) it has, divided by the number of interpretations that all possible grammatically correct sentences have. The refinement of that equation to equation 6 now says that sentences with more meaning are less probable, because they are compatible with fewer interpretations, and that for any given degree of meaningfulness, sentences with a higher semantic potential are more probable. (That is, precise sentences are harder to write, but it’s easier to write a precise sentence with more words, and it’s harder to pack all your meaning into just a few very precise words.)

3.4. From Ensembles to Vectors

Most word embedding models generate word vectors by using a supervised or semi-supervised model to predict cooccurrences, and the vectors themselves aren’t significant outside that predictive frame. But the picture is quite different for statistical-mechanical models such as this one. One of the most elegant properties of partition functions is that a wide range of thermodynamic quantities can be expressed directly as partial derivatives of the partition function or its logarithm.

For example, suppose we would like to determine the number of particles of a particular kind present in all possible states of our system (N_k), and take the average. We can calculate that value by taking the partial derivative of the logarithm of equation 6 with respect to the chemical potential of that species, and dividing out $\beta = \frac{1}{k_B T}$.

$$\langle N_k \rangle = \frac{1}{\beta} \frac{\partial \ln Z(\mu_k)}{\partial \mu_k} \quad (7)$$

Since $\frac{\partial \ln f(x)}{\partial x} = \frac{\partial f(x)}{\partial x} / f(x)$, this simplifies to a probability-weighted sum of N_k counts divided by β , effectively an arbitrary constant multiplier. Shifting it to the left hand side of the equation gives

$$\beta \langle N_k \rangle = \frac{\frac{\partial}{\partial \mu_k} Z(\mu_k)}{Z(\mu_k)} = \sum_i N_{k,i} \frac{e^{\beta(\mu_1 N_{1,i} + \mu_2 N_{2,i} + \dots + \mu_k N_{k,i} - E_i)}}{Z} = \sum_i N_{k,i} p_i \quad (8)$$

This line of reasoning can be extended to second partial derivatives. The variance of N_k is given by

$$\beta [\langle N_k^2 \rangle - \langle N_k \rangle^2] = \frac{\partial^2 \ln Z(\mu_k)}{\partial \mu_k^2} \quad (9)$$

Similarly, the covariance of N_k and N_j is a mixed partial derivative.

$$\beta [\langle N_k N_j \rangle - \langle N_k \rangle \langle N_j \rangle] = \frac{\partial^2 \ln Z(\mu_k, \mu_j)}{\partial \mu_k \mu_j} \quad (10)$$

These last two equations can be used to construct a matrix that has two simultaneous meanings. It is, first, a covariance matrix that describes the way particle counts are correlated with one another in the system. But it is also a Hessian matrix of second partial derivatives, meaning that it describes the way small modifications to the chemical potential terms change the overall partition function, shifting its energy balance across all possible system microstates. This means that even if we can't construct the partition function itself, we can in principle measure the covariance of particles empirically, and use the resulting matrix to reconstruct information about the partition function and the thermodynamic ensemble it describes.

If we translate this into linguistic terms, we find that by taking empirical measurements of word cooccurrence, we are also constructing the Hessian of a linguistic partition function that describes how changes to the meaning of one word affect the meaning of another. The columns of that matrix are word vectors. When two columns are similar, small modifications to the meanings of the associated words have similar effects on the language as a whole. That is what it means, in the context of this model, for two words to be similar. Line integrals through the Hessian field in a given neighborhood can also be approximated by adding and subtracting these vectors, giving a more precise interpretation to the formulas used to represent analogies. Analogies are valid when they correspond to two different line integrals through a conservative Hessian tensor field, beginning at the same point and ending close to the same point, and therefore having similar final values.

3.5. Implementation

Constructing a practical implementation of this model requires that we determine the values for two sets of parameters: the potential for each word in the vocabulary, and the energy level for each sentence. The simplest approach to this problem is to set all potential terms to zero, and all energy terms to one. The covariance matrix that results from these choices is identical to the one given by directly counting word cooccurrences. Alternative schemes will change the weights given to each of the sentences, yielding a modified covariance matrix that is likely to give better meaning representations. For performance reasons, some form of dimension reduction is also necessary, but has no theoretical significance at all. In practice, random projection (as in [25]) works well, especially after implementing some of the preprocessing and hyperparameter selection recommendations in [16], which may be compensating for the deficiencies that result from setting the energy and potential terms to constants.

The problem of selecting energy and potential terms in a more principled way is left to other work. But it is worth considering briefly, since it illustrates some interesting properties of the model. First, in this model, the same sequence of words could appear twice with different energy levels, and therefore different probabilities, depending on context. Second, there may be a way to make predictions that link the semantic potential of given terms to known lexicographical properties of those terms, such as the approximate number of senses the word has. And finally, the partition function described here is not the only possible partition function that might be applied to language. Partition functions based on word pairs, sequences, or even attention mechanisms could be used to model language within this framework, all broadly interpretable in the same way.

4. Discussion

Few of the ideas presented here are new. The fact that word vectors contain distributional information that allows them to measure word similarity has been known for decades. Ideas from statistical mechanics have been applied to language modeling, machine learning, and information retrieval problems for decades. And for the last few years, there has been a steady stream of work demonstrating that language model X reduces to language model Y in some limit. But none of this has shown how these models could capture information about interpretation or meaning. Trained on linguistic form alone, these models have no evidence showing how linguistic forms map to mental models, concepts, narratives, or any other representations of things outside of language.

The claim that statistical models can infer things about meaning from linguistic form alone thus faces a high burden of proof. And while there has been a proliferation of models that do appear to support that claim, they all work on slightly different principles, and produce slightly different results. This undercuts attempts at meta-induction; many small bodies of evidence based on different principles of operation do not add up to one large body of evidence. And so justified skepticism remains.

What is new about the model proposed here is that it is general enough to explain the success of many of these models without reference to the details of their operation. Fundamentally, any model that is able to predict linguistic cooccurrences can be reinterpreted as an implicit partition function along the lines proposed here. So reinterpreted, we can argue that distributional information about language is linked by a precise mathematical structure to specific facts about how words signify. Those facts are limited; they do not include any information about *what* words, sentences, or longer fragments of language talk about. But they do include information about *how many* interpretations might be applied to those units of language, and how those interpretations correlate with one another at a macroscopic level.

What unites all of these models, under this theory, is that they effectively assume that meaning, quantified appropriately, is conserved, and that units of language—be they letters, words, n-grams, or longer phrases—are also conserved. It's not yet clear what these assumptions might mean in linguistic terms, but they are crucial to the derivation of a partition function that can relate the statistics of linguistic form to an unknown reservoir of meaning.

These models must also make a third assumption: language exists in a state of equilibrium with its reservoir of meaning. That assumption is unlikely to hold in general. If this way of thinking about language modeling is sound, then an important project will be to understand when the assumption of equilibrium is justified, and when it is not. It's likely that during periods of rapid linguistic change, for example, the equilibrium assumption will not be valid. In that case, methods that can model far-from-equilibrium systems will be required. Since non-equilibrium thermodynamics is a field still in its infancy [6], there will be much work to be done, and many tasks that remain impossible without domain expertise. Nonetheless, a deeper understanding of the meaning of these assumptions promises to clarify when and how language models can infer meaning from linguistic form alone.

References

- [1] M. Antoniak and D. Mimno. “Evaluating the Stability of Embedding-based Word Similarities”. In: *TACL* 6 (2018), pp. 107–119. URL: <https://transacl.org/ojs/index.php/tacl/article/view/1202>.
- [2] S. Arora et al. “A Latent Variable Model Approach to PMI-based Word Embeddings”. In: *TACL* 4 (2016), pp. 385–399. DOI: 10.1162/tacl_a_00106. URL: <https://www.aclweb.org/anthology/Q16-1028>.
- [3] J. Baez and M. Stay. “Algorithmic thermodynamics”. In: *Mathematical Structures in Computer Science* 22.5 (2012), pp. 771–787. DOI: 10.1017/S0960129511000521.
- [4] E. M. Bender and A. Koller. “Climbing towards NLU: On Meaning, Form, and Understanding in the Age of Data”. In: *ACL 2020*. Online: Association for Computational Linguistics, July 2020, pp. 5185–5198. URL: <https://www.aclweb.org/anthology/2020.acl-main.463>.
- [5] T. B. Brown et al. “Language Models are Few-Shot Learners”. In: *CoRR* abs/2005.14165 (2020). URL: <https://arxiv.org/abs/2005.14165>.
- [6] J. England. “Dissipative adaptation in driven self-assembly”. In: *Nature Nanotechnology* 10 (Nov. 2015), pp. 919–923. DOI: 10.1038/nnano.2015.250.
- [7] J. Firth. “A synopsis of linguistic theory 1930–55”. In: *Studies in linguistic analysis. The Philological Society, Oxford* (1957), pp. 1–32.
- [8] M. Gavin. “Vector Semantics, William Empson, and the Study of Ambiguity”. In: *Critical Inquiry* 44.4 (2018), pp. 641–673. DOI: 10.1086/698174.
- [9] A. Gladkova, A. Drozd, and S. Matsuoka. “Analogy-based detection of morphological and semantic relations with word embeddings: what works and what doesn’t”. In: *SRW@HLT-NAACL 2016, San Diego California, USA, June 12-17, 2016*. The Association for Computational Linguistics, 2016, pp. 8–15. DOI: 10.18653/v1/n16-2002. URL: <https://doi.org/10.18653/v1/n16-2002>.
- [10] W. L. Hamilton, J. Leskovec, and D. Jurafsky. “Diachronic Word Embeddings Reveal Statistical Laws of Semantic Change”. In: *ACL 2016*. Berlin, Germany: Association for Computational Linguistics, Aug. 2016, pp. 1489–1501. DOI: 10.18653/v1/P16-1141. URL: <https://www.aclweb.org/anthology/P16-1141>.
- [11] Z. S. Harris. “Distributional structure”. In: *Word* 10.2-3 (1954), pp. 146–162.
- [12] R. J. Heuser. “Word Vectors in the Eighteenth Century”. In: *DH 2017, Montréal, Canada, August 8-11, 2017, Conference Abstracts*. Ed. by R. Lewis et al. Alliance of Digital Humanities Organizations (ADHO), 2017, pp. 256–259. URL: <https://dh2017.adho.org/abstracts/582/582.pdf>.
- [13] G. Jawahar, B. Sagot, and D. Seddah. “What Does BERT Learn about the Structure of Language?”. In: *ACL 2019*. Florence, Italy: Association for Computational Linguistics, July 2019, pp. 3651–3657. DOI: 10.18653/v1/P19-1356. URL: <https://www.aclweb.org/anthology/P19-1356>.
- [14] A. Kolchinsky and D. H. Wolpert. “Semantic information, autonomous agency and non-equilibrium statistical physics”. In: *Interface Focus* 8.6 (2018), p. 20180041.

- [15] O. Levy and Y. Goldberg. “Neural Word Embedding as Implicit Matrix Factorization”. In: *NIPS 2014, December 8-13 2014, Montreal, Quebec, Canada*. 2014, pp. 2177–2185. URL: <http://papers.nips.cc/paper/5477-neural-word-embedding-as-implicit-matrix-factorization>.
- [16] O. Levy, Y. Goldberg, and I. Dagan. “Improving Distributional Similarity with Lessons Learned from Word Embeddings”. In: *TACL 3 (2015)*, pp. 211–225. DOI: 10.1162/tacl_a_00134. URL: <https://www.aclweb.org/anthology/Q15-1016>.
- [17] M. Masterman. “Semantic algorithms”. In: *Language, Cohesion and Form*. Studies in Natural Language Processing. Cambridge University Press, 2005, pp. 253–280. DOI: 10.1017/CBO9780511486609.012.
- [18] M. Masterman (Braithwaite). “XI.—Words”. In: *Proceedings of the Aristotelian Society* 54.1 (July 2015), pp. 209–232. ISSN: 0066-7374. DOI: 10.1093/aristotelian/54.1.209. URL: <https://academic.oup.com/aristotelian/article-pdf/54/1/209/5256573/aristotelian54-0209.pdf>.
- [19] T. McCoy, E. Pavlick, and T. Linzen. “Right for the Wrong Reasons: Diagnosing Syntactic Heuristics in Natural Language Inference”. In: *ACL 2019*. Florence, Italy: Association for Computational Linguistics, July 2019, pp. 3428–3448. DOI: 10.18653/v1/P19-1334. URL: <https://www.aclweb.org/anthology/P19-1334>.
- [20] T. Mikolov et al. “Efficient Estimation of Word Representations in Vector Space”. In: *ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*. Ed. by Y. Bengio and Y. LeCun. 2013. URL: <http://arxiv.org/abs/1301.3781>.
- [21] D. M. Mimno and L. Thompson. “The strange geometry of skip-gram with negative sampling”. In: *EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*. 2017, pp. 2873–2878. URL: <https://aclanthology.info/papers/D17-1308/d17-1308>.
- [22] T. Niven and H.-Y. Kao. “Probing Neural Network Comprehension of Natural Language Arguments”. In: *ACL 2019*. Florence, Italy: Association for Computational Linguistics, July 2019, pp. 4658–4664. DOI: 10.18653/v1/P19-1459. URL: <https://www.aclweb.org/anthology/P19-1459>.
- [23] J. Pennington, R. Socher, and C. D. Manning. “Glove: Global Vectors for Word Representation”. In: *EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*. 2014, pp. 1532–1543. URL: <http://aclweb.org/anthology/D/D14/D14-1162.pdf>.
- [24] M. Peters et al. “Deep Contextualized Word Representations”. In: *HLT-NAACL 2018*. New Orleans, Louisiana: Association for Computational Linguistics, June 2018, pp. 2227–2237. DOI: 10.18653/v1/N18-1202. URL: <https://www.aclweb.org/anthology/N18-1202>.
- [25] B. Schmidt. “Stable Random Projection: Lightweight, General-Purpose Dimensionality Reduction for Digitized Libraries”. In: *Journal of Cultural Analytics* (2018). DOI: 10.22148/16.025. URL: <https://culturalanalytics.org/article/11033>.
- [26] J. R. Searle. “Minds, Brains, and Programs”. In: *Behavioral and Brain Sciences* 3.3 (1980), pp. 417–57. DOI: 10.1017/s0140525x00005756.

- [27] N. Srivastava, R. Salakhutdinov, and G. E. Hinton. “Modeling Documents with Deep Boltzmann Machines”. In: *Proceedings of the Twenty-Ninth Conference on Uncertainty in Artificial Intelligence, UAI 2013, Bellevue, WA, USA, August 11-15, 2013*. Ed. by A. Nicholson and P. Smyth. AUAI Press, 2013. URL: https://dslpitt.org/uai/displayArticleDetails.jsp?mmnu=1%5C&smnu=2%5C&article%5C_id=2423%5C&proceeding%5C_id=29.
- [28] G. J. Stephens and W. Bialek. “Statistical mechanics of letters in words”. In: *Physical Review E* 81.6 (June 2010). ISSN: 1550-2376. DOI: 10.1103/physreve.81.066119. URL: <http://dx.doi.org/10.1103/PhysRevE.81.066119>.
- [29] A. Vaswani et al. “Attention Is All You Need”. In: *CoRR* abs/1706.03762 (2017). URL: <http://arxiv.org/abs/1706.03762>.
- [30] W. Weaver. “Translation”. In: *Machine translation of languages: fourteen essays*. MIT and Wiley, 1955, pp. 15–23.