

On the Information Content of Predictions in Word Analogy Tests

Jugurta Montalvão

Abstract—An approach is proposed to quantify, in bits of information, the actual relevance of analogies in analogy tests. The main component of this approach is a *soft* accuracy estimator that also yields entropy estimates with compensated biases. Experimental results obtained with pre-trained GloVe 300-D vectors and two public analogy test sets show that proximity hints are much more relevant than analogies in analogy tests, from an information content perspective. Accordingly, a simple word embedding model is used to predict that analogies carry about one bit of information, which is experimentally corroborated.

Index Terms—Word embeddings, Word analogy, GloVe, Information content, Soft accuracy.

I. INTRODUCTION

TEXTS in natural language can be regarded as a kind of signal, but its non-numerical nature obviously prevents the use of most signal processing methods. As for signal prediction, however, the coding of words as continuous real-valued vectors seems to nicely circumvent this limitation. Indeed, Mikolov et al. [1] proposed a kind of prediction task where a machine was expected to solve analogies, thus predictions, such as the automatic guessing that the word like “king” is similar to “man” in the same sense as “queen” is similar to “woman.” Accordingly, it is expected that a good word embedding for words in English would associate vectors to words, say, \vec{a} , \vec{b} , $\vec{\alpha}$ and $\vec{\beta}$, for “king”, “queen”, “man” and “woman”, respectively, in such a geometric configuration that the following approximation should hold:

$$\vec{a} - \vec{\alpha} + \vec{\beta} \approx \vec{b}. \quad (1)$$

This is referred to as the *vector offset method* (VOM) for solving word analogies, which has become a standard tool for semantic models in vector-spaces.

Since [1] was published, word embeddings obtained with artificial neural networks became popular in applications ranging from document representation [2] to sentiment analysis [3]. Besides, word analogy became one of the tests frequently used to evaluate word embeddings (e.g. [4], [5], [6], [7], [8]), in spite of the VOM been criticized in its original formulation and results, to the point of suspicion regarding the actual relevance of analogies in the high accuracies frequently obtained. Remarkably, even the authors of [1] found “somewhat surprisingly” that analogy questions could be answered via the VOM, and to experimentally probe its prediction power, they prepared a list of 19,544 prediction tests, split into 8,869

semantic and 10,675 syntactic ones [9], henceforth referred to as the Google Analogy Test Set (GATS). All tests followed the same format as w_a is to w_b as w_α is to w_β , where w_x stands for the word associated to vector \vec{x} .

Experiments done with the GATS yielded impressive results in terms of prediction accuracies with VOM, but it also attracted some criticism, ranging from its statistical consistency [10] and foundations [11] to its optimization details [12], to the point that even the validity of the arithmetic analogy test has been questioned [13]. For instance, as pointed out in [10], frustrating experimental results in analogy tests are observed if the instruction to discard input question words (before starting the output word search) is disregarded.

This instruction is explicitly given in [1], and the corresponding frustrating results suggest that vector

$$\vec{\phi}_{\alpha,\beta} = \vec{\beta} - \vec{\alpha}, \quad (2)$$

which would encode the most valuable part of the analogy between pairs of words, may not play this expected role. On the other hand, proximity between similar words in the embedding is also expected to play a relevant role, as studied by Lund and Burgess [14], long before VOM was proposed. According to [14], neighbourhood surroundings are akin to a *semantic fields* by themselves, where similar representations tend to cluster words that can be substituted by each other in a context.

In [15], the relevance of vector $\vec{\phi}_{\alpha,\beta}$ for linguistic relations in word embeddings was studied through the pairing consistency score between $\vec{\phi}_{\alpha,\beta}$ and

$$\vec{\phi}_{a,b} = \vec{b} - \vec{a}, \quad (3)$$

where some pairs in analogy test sets were randomly shuffled, and Area Under the Curve values were subsequently computed to compare angle distributions with and without pair shuffling.

Alternatively, the VOM formulation may be regarded as a word guessing game where w_b is to be guessed from two separate hints, namely: that w_b is similar to w_a , and that this similarity is analogous to that between w_α and w_β . Because words are represented as real-valued vectors, this guessing game is also tantamount to a signal prediction where \vec{a} plays the role of current signal instance, and $\vec{\phi}_{\alpha,\beta}$ should be a noisy direction for estimating \vec{b} as a quantized version of prediction $\vec{p} = \vec{a} + \vec{\phi}_{\alpha,\beta}$. This signal prediction perspective further suggests that a simpler zero-order predictor would be $\vec{p}_1 = \vec{a}$, in which $\vec{\phi}_{\alpha,\beta}$ is not taken into account at all. This simpler predictor corresponds to the use of the first hint only, in the word guessing game, whereas the full predictor $\vec{p}_2 = \vec{a} + \vec{\phi}_{\alpha,\beta}$ uses the two hints.

J. Montalvão is with the Department of Electrical Engineering, Federal University of Sergipe, São Cristóvão, 49000-000.

E-mail: jmontalvao@academico.ufs.br

Digital Object Identifier: 10.14209/jcis.2022.18

In this work, we take this guessing game perspective to measure information content in popular analogy tests. In Section II, the proposed approach is explained in terms of usual accuracy, whereas in Section III an alternative accuracy measurement is introduced as a better estimator, for the purpose of this work. Experimental results are presented in Section IV, thus leading to the formulation of an analytical model for analogies in sparse word embeddings, in Section V. Conclusions are presented in Section VI.

II. MEASURING INFORMATION IN ANALOGY TESTS

From an information perspective, each analogy test may be regarded as a guessing game where w_b is a target symbol (word) to be guessed in a finite set U of M symbols. This perspective is illustrated in Fig. 1, whereas Fig. 2 illustrates a geometrical perspective of hints in analogy tests. Accordingly, two hints are provided as:

- h1** Similarity: w_b is similar to w_a . Therefore, in the word embedding \vec{b} is likely to be a near-neighbour of $\vec{p}_1 = \vec{a}$;
- h2** Analogy: w_α relates to w_β in the same manner as w_a relates to w_b , then \vec{b} is likely to be a near neighbour of $\vec{p}_2 = \vec{a} + \vec{\beta} - \vec{\alpha}$.

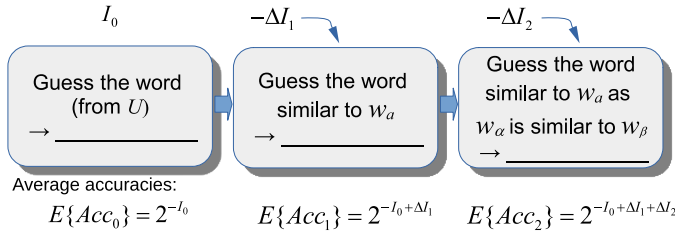


Fig. 1. Word guessing in three levels of difficulty. From left to right: (1) Randomly guess a word in a set U of M words; (2) Guess a word w_b whose corresponding vector is around \vec{a} (3) Guess a word w_b whose corresponding vector is around \vec{a} , toward $\vec{\phi}_{\alpha,\beta}$.

For the sake of a more intuitive presentation of ideas, the definition found in [16] for the Shannon information content of an outcome x is used here, as:

$$I(x) = \log_2 \frac{1}{\Pr(X = x)} \quad (4)$$

which turns out to be, in Shannon's original work [17], the proposed "measure of how much *choice* is involved in the selection of the event or of how uncertain we are of the outcome". The *information content* used in this work is a short for the *Shannon information content*, which is also defined for a random categorical outcome, such as the choice of a word. Therefore, Shannon's entropy is defined as the average information content. These connections to the Shannon's seminal work are further discussed in the Appendix.

Notice that, as in [18], accuracies can be alternatively thought in terms of entropy, through the concept of *effective cardinality*, as illustrated in the Appendix. Accordingly, any classification/detection problem, with a given accuracy, A , is analogous to another problem of finding a single target element in a *chimeric* set of $C = 1/A$ equally likely ones, where C is the effective cardinality of the *chimeric* set, and

$H = \log_2(C)$ is the corresponding entropy, or its average information content, in bits. Effective cardinality is used here because it is assumed to intuitively reflect difficulty levels of a guessing game.

In the word guessing game illustrated in Fig. 1, hints can be regarded as information content injections, $\Delta I_1 = I_0 - I_1$ and $\Delta I_2 = I_1 - I_2$, in bits, where only ΔI_2 corresponds to the piece of information that gives name to (analogy) tests.

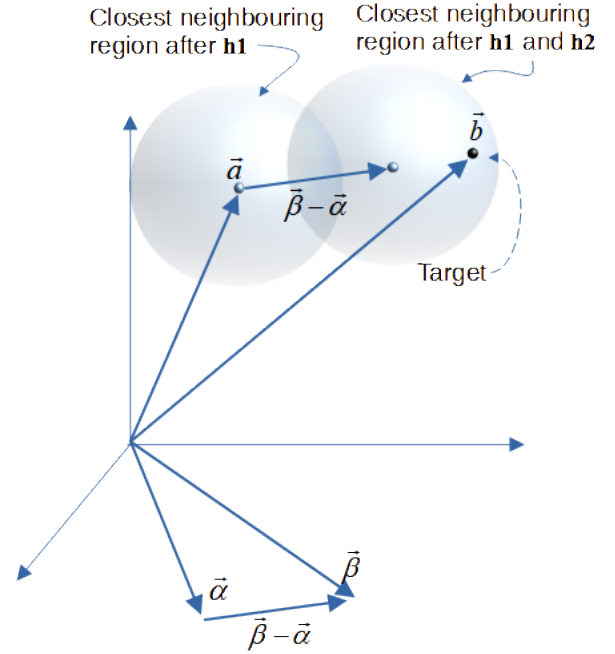


Fig. 2. Illustration of how the two hints, **h1** and **h2**, modify the neighbouring region where the target vector is searched. Ideally, a good hint **h2** is expected to move the search region toward the target, b , which is associated to the correct answer w_b .

Before any hint is provided, guessing the right word is worth I_0 bits. In natural language corpus, words are far from equally distributed, but there is no clear reason for assuming that a target word arbitrarily chosen for guessing purposes should follow any non-uniform distribution. Therefore, a conservative approach shall assume that $I_0 = \log_2(M)$ bits, where M is the cardinality of the set of all symbols, U , from which w_b is randomly drawn. For instance, for a vocabulary of $M = 400,000$ unique words in English, the guessing of w_b without any hint has an information content of about 18.6 bits. Hint **h1**, when provided, may have the power to shrink the search space to the near-neighbours of w_a , regardless the direction of these neighbours, which reduces the amount of uncertainty about w_b to $I_0 - \Delta I_1$ bits. Eventually, **h2** may further reduce the uncertainty about w_b by moving the neighbourhood search in the direction $\vec{\phi}_{\alpha,\beta}$. This further reduction is expected even for a weak alignment between $\vec{\phi}_{\alpha,\beta}$ and $\vec{\phi}_{a,b}$, as studied in [15], and finally the amount of uncertainty about w_b is reduced to $I_0 - \Delta I_1 - \Delta I_2$ bits.

Entropy measures reflect uncertainties about w_b , and are theoretically related to accuracies [18]. Indeed, average accuracies

can be expressed as

$$E\{Acc_1\} = \frac{1}{2^{I_0 - \Delta I_1}} \quad (5)$$

and

$$E\{Acc_2\} = \frac{1}{2^{I_0 - \Delta I_1 - \Delta I_2}}, \quad (6)$$

where Acc_1 and Acc_2 are random variables that model experimental accuracy instances, acc_1 and acc_2 , of the guessing game, after $\mathbf{h1}$ and $\mathbf{h2}$ are provided, respectively.

In turn, experimental accuracies, acc_1 and acc_2 , can be used to estimate the amount of information carried out by each hint, as:

$$\widehat{\Delta I}_1 = I_0 - \log_2(1/acc_1), \quad acc_1 \neq 0 \quad (7)$$

and

$$\widehat{\Delta I}_2 = \log_2(1/acc_1) - \log_2(1/acc_2), \quad acc_2 \neq 0 \quad (8)$$

III. SOFT ACCURACY

For finite sequences of a guessing game, accuracy estimates are typically spiky and slowly convergent, due to the binary nature of instances. To circumvent these problems, an improved accuracy estimator is proposed in this section, referred to as the *soft accuracy*, where instead of a (binary) hard decision/guess per experiment run, an effective cardinality is estimated from the position of the right answer. As illustrated in Fig. 3, the instance estimate for effective cardinality, c , in a single experiment, is defined as

$$c = 2o_t - 1 \quad (9)$$

where o_t is the sorted position of the target in ascending order of distance (or descending order or similarity), with regard to the hinted initial search position, i.e. prediction \vec{p} . In other words, o_t stands for an instance of the random variable O_t , which models the sorted position of the target, thus it takes values from $\{1, 2, \dots, M\}$.

This formulation stems from the belief that, in a huge set of possible answers, a small subset of c answers around a given prediction are (almost) equally likely to be the target. Therefore, it is assumed that every observed instance of o_t relates to a corresponding unobserved instance of c according to $o_t = (c + 1)/2$.

Therefore, c represents an instance of the random variable $C = 2O_t - 1$, that models the effective cardinality of a search space around predictions, and takes values from $\{1, 3, 5, \dots, 2M - 1\}$. Accordingly, each independent instance of c assumes that o_t is one out of $2o_t - 1$ neighbours of \vec{p} that are equally likely to be the target, thus c roughly represents an average number of points around \vec{p} . In logarithmic scale, every independent guess, g , yields the following instance of information content of the right guess outcome (a hit):

$$\hat{h}_g = \log_2(c) \quad (10)$$

which is also the entropy of a hypothetical random variable uniformly distributed among c values. Thus, over G independent runs of the game, a “naive” entropy (i.e. average information content) estimate would be approximated as

$$\hat{H}_G \approx \frac{\sum_{g=1}^G \hat{h}_g}{G} \quad (11)$$

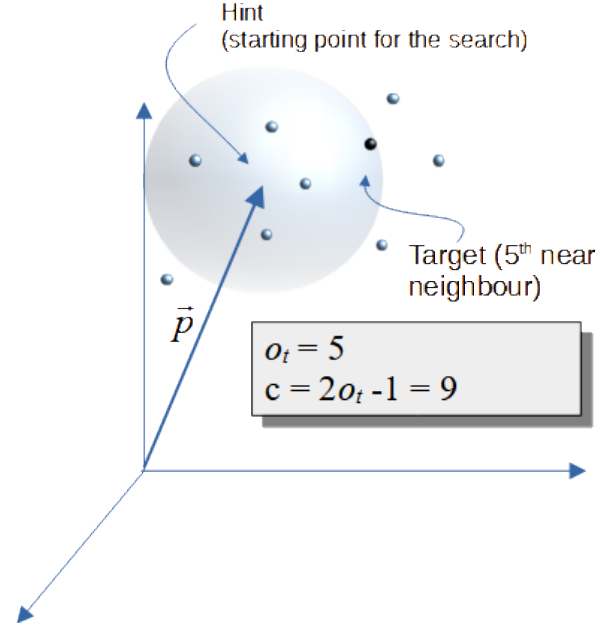


Fig. 3. Illustration of the effective cardinality estimate for a single run of the guessing game, where the hint is that the target is around the point indicated by vector \vec{p} . However, the target is actually the 5th near neighbour from that point, thus $o_t = 5$, and $c = 9$ indicates that the target is the median of a set of 9 neighbours, taken as single independent instance of the search space cardinality.

Alas, this kind of entropy estimator has long been known to have a bias induced by the logarithm in Eq. 10. Probably the simplest compensation for this kind of bias was proposed by G. Miller [19], in 1955. However, the entropy estimator in Eq. 11 is not exactly the same studied by G. Miller, and a specific bias analysis is necessary here, where what we want is an estimate of $\log_2(E\{C\})$, but we estimate $E\{\log_2(C)\}$ instead. The above mentioned estimator bias is the resulting difference, B , formulated as

$$E\{\log_2(C)\} = \log_2(E\{C\}) + B \quad (12)$$

According to Eq. 9, for target positions limited to N near neighbours, the left side of Eq. 12 can be expanded as

$$E\{\log_2(C)\} = \sum_{n=1}^N \Pr(c_n) \log_2(2n - 1). \quad (13)$$

If we further assume that all N neighbours are equally likely to be the target, then it follows that

$$E\{\log_2(C)\} = \frac{1}{N} \sum_{n=1}^N \log_2(2n - 1). \quad (14)$$

By adding and subtracting $\log_2(N)$ in the right side of this equation, and after a few straightforward algebraic manipulation, we obtain

$$E\{\log_2(C)\} = \log_2(N) + \frac{1}{N} \sum_{n=1}^N \log_2 \frac{2n - 1}{N} \quad (15)$$

On the other hand, for a uniform distribution, we have

$$E\{C\} = \frac{1}{N} \sum_{n=1}^N (2n - 1) = N, \quad (16)$$

and from Eq. 15 and Eq. 16, we obtain

$$E\{\log_2(C)\} = \log_2(E\{C\}) + \frac{1}{N} \sum_{n=1}^N \log_2 \frac{2n-1}{N} \quad (17)$$

Finally, by comparing Eq. 17 to Eq. 12 we conclude that

$$B = \frac{1}{N} \sum_{n=1}^N \log_2 \frac{2n-1}{N} \quad (18)$$

Therefore, given N , and under the assumption that all N neighbours are equally likely to be the target, a compensated entropy estimate is given by

$$\log_2(E\{C\}) = \hat{H}_G - B \quad (19)$$

and the soft accuracy can be obtained as

$$acc = 2^{-\hat{H}_G + B}. \quad (20)$$

An illustrative comparison between “hard” and soft accuracies is presented in Fig. 4, for a guessing game with $N = 10$ and where a target, o_t , and a hint, o_{hint} , are randomly chosen in each independent run of the game. After G such runs, the hard accuracy is given by the number of times o_{hint} coincides with o_t divided by G , whereas the soft accuracy is obtained as in Eq. 20. Both accuracy estimates converge to $1/N = 0.1$, but the convergence of the soft one is faster and less susceptible to wide deviations. Therefore, in this work, rather than accuracy itself, entropy estimates as provided in Eq. 19 are to be used.

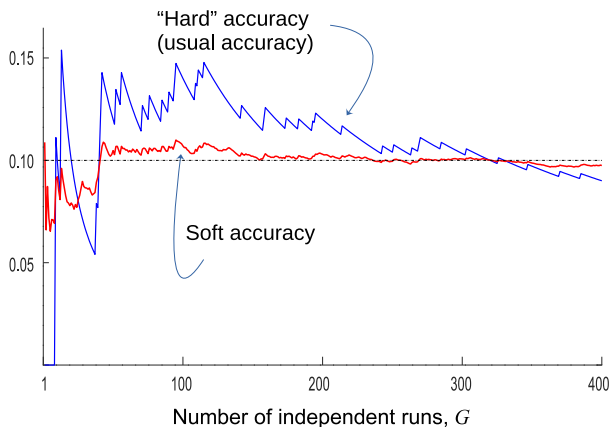


Fig. 4. Illustration comparison between hard and soft accuracies for a guessing game with $N = 10$, where hints are completely random (pseudo-hints). Both accuracies converge to $1/N = 0.1$.

A possible drawback of this proposed soft accuracy would be that N (the effective number of neighbours) must be known *a priori* to allow a proper bias compensation. Helpfully, the range of values that B assumes is small. For instance, for N in the range from 2 to 10,000, B varies from about -0.20 to about -0.44 . In the experimental part of this work, we take advantage of it to assume a single heuristic value for (unknown) N , and then use the corresponding B throughout all experiments. Indeed, in analogy tests, the target word is frequently either the first or the second near neighbour, sometimes the third, thus suggesting that the average number

of near neighbours of predictions that are equally likely to be the target is not far from 2. Therefore, in experiments done in Section IV, the bias compensation value was arbitrarily set to -0.25 .

IV. EXPERIMENTS

In what follows, Equations 7 and 8, along with the soft accuracy defined in Eq. 20, are articulated to build a pragmatic approach to measure the relevance of each hint that composes an analogy test, through the following kinds of experiments:

- **Single-hint experiment:** just word w_a is taken into account, and the nearest neighbour of prediction $\vec{p}_1 = \vec{a}$, say \vec{g}_1 , is taken as a tentative guess for \vec{b} . All vectors are normalized and proximity scores are obtained as inner products,

$$s_m = \langle \vec{p}_1, \vec{v}_m \rangle, \vec{v}_m \neq \vec{a}, \quad (21)$$

where \vec{v}_m is drawn from a set of $M - 1$ vectors. In a test set with G questions, for each question where $o_b = o_t$ is the order of the target score s_b , obtained for \vec{v}_b , c_g is obtained according to Eq. 9, representing the effective cardinality instance associated to the g^{th} question/guess. Then an estimate \hat{H}_G is accumulated as in Eq. 11, thus allowing the computation of a soft accuracy instance, acc_1 , as in Eq. 20.

- **Two-hints experiment:** a prediction is given by $\vec{p}_2 = \vec{a} + \vec{\beta} - \vec{\alpha}$. Again, all vectors are normalized and proximity scores are obtained as

$$s_m = \langle \vec{p}_2, \vec{v}_m \rangle, \vec{v}_m \notin \{\vec{a}, \vec{\alpha}, \vec{\beta}\}, \quad (22)$$

where \vec{v}_m is now drawn from a set of $M - 3$ vectors, and the corresponding soft accuracy instance, acc_2 , is obtained.

The first set of experiments was performed with a pre-trained set of $M = 400,000$ words in English, encoded with GloVe [20] as 300-D vectors. The corresponding embedding of this publicly available dataset is referred to as the *Wikipedia 2014 + Gigaword 5*. As for the tests, we used both the GATS [9], in Table I, and the Bigger Analogy Test Set (BATS)[21], in Table II.

As expected, all results are consistent with the conclusions already published in many former works (e.g. [10], [11], [12], [13] and [15]), that the analogy itself is not as important as accuracies based on both hints may suggest. Additionally, presented values in bits for $\widehat{\Delta I}_1$ and $\widehat{\Delta I}_2$ provide a suitable quantitative perception of it. For instance, while in Table I the GATS have analogies with average information content of about 1.8 bits, which is already small, as compared to the average 15.3 bits for $\widehat{\Delta I}_1$, results in Table II corroborate the expectation that BATS is a more challenging set of analogies, yielding even a loss of information content of about 1 bit, for the *Lexicographic semant.* subset. This loss is expressed as a negative value that means that **h2** is a disturbing noise — rather than a true hint — that doubles the average search space for target words. It is worth noticing that this negative value is not an information content, which cannot be negative, by definition. A simple illustration of this kind of negative result is presented in the Appendix.

TABLE I
SOFT ACCURACIES (IN %) AND INFORMATION CONTENT VARIATIONS (IN BITS) FOR THE GATS.

Subsets	Soft accuracy (single-hint)	Soft accuracy (two-hints)	$\widehat{\Delta I}_1$	$\widehat{\Delta I}_2$
capital-common-countries	25.5	80.2	16.6	1.7
capital-world	27.5	81.0	16.7	1.6
city-in-state	4.5	28.2	14.1	2.6
currency	0.1	0.5	7.9	3.0
family	30.1	68.5	16.9	1.2
gram1-adjective-to-adverb	3.2	5.2	13.6	0.7
gram2-opposite	1.0	2.5	12.0	1.3
gram3-comparative	17.4	68.9	16.1	2.0
gram4-superlative	1.3	46.6	12.4	5.1
gram5-present-participle	27.6	44.6	16.8	0.7
gram6-nationality-adjective	29.5	68.8	16.8	1.2
gram7-past-tense	21.6	37.5	16.4	0.8
gram8-plural	42.4	59.5	17.4	0.5
gram9-plural-verbs	7.3	36.1	14.8	2.3

TABLE II
SOFT ACCURACIES (IN %) AND INFORMATION CONTENT VARIATIONS (IN BITS) FOR THE BATS.

Subsets	Soft accuracy (single-hint)	Soft accuracy (two-hints)	$\widehat{\Delta I}_1$	$\widehat{\Delta I}_2$
Inflectional morph.	7.5	27.3	14.9	1.9
Derivational morph.	0.2	0.4	9.6	0.9
Encyclopedic semant.	0.4	2.0	10.6	2.4
Lexicographic semant.	0.7	0.3	11.4	-1.0

V. SIMPLE MODEL FOR ANALOGIES IN WORD EMBEDDINGS

Experimental values of $\widehat{\Delta I}_2$ in both tables I and II yield an average value of about 1.6 bits, which is much less than any value of $\widehat{\Delta I}_1$. To afford some understanding of such disparity, even a coarse model for the word embedding can be of help, as follows: let all M words to appear in pairs (\vec{a}_m, \vec{b}_m) , $m = 1, 2, \dots, M/2$, so that they almost satisfy $\vec{b}_m \approx \vec{a}_m + \vec{\phi}_{\alpha, \beta}$, where a constant $\vec{\phi}_{\alpha, \beta}$ encodes a single unique analogy relationship for all pairs of word in this coarse model. Besides, the embedding is assumed to be sparse (as actual word embeddings [15]), therefore most directions point out toward space regions where points representing words are far apart from each other.

Before any hint is provided, an observer cannot know whether a word is of kind a (start-word) or b (end-word). When only hint $\mathbf{h2}$ is provided, direction $\vec{\phi}_{\alpha, \beta}$ is given, but the start-word is not. Consequently, $\vec{\phi}_{\alpha, \beta}$ must be added to all words, as illustrated in Fig. 5, because the guesser is unaware of which are the start-words, thus yielding M predictions. However, only for the $M/2$ actual start-words the approximation

$$\vec{b}_m \approx \vec{a}_m + \vec{\phi}_{\alpha, \beta} \quad (23)$$

is warranted, which is likely to place the actual target among the $M/2$ actual end-words, whereas the embedding sparseness

is likely to lead most predictions from end-words,

$$\vec{p} = \vec{b}_m + \vec{\phi}_{\alpha, \beta}, \quad (24)$$

to regions of the embedding where even the first near neighbour is far away. Therefore one should expect to find the target \vec{b}_t among the $M/2$ closest near neighbours (NN) of all M predictions.

Again, due to the approximation

$$\vec{b}_m \approx \vec{a}_m + \vec{\phi}_{\alpha, \beta}, \quad (25)$$

the $M/2$ closest near neighbours are expected to be almost equally likely to be the tentative guess. According to equations from Eq. 16 to 20, if we replace the N equally probable symbols with $M/2$, the expected soft accuracy is indeed $2/M$, for the search space is reduced, on average, to $M/2$ (almost) equally likely predictions, and

$$I_0 - \log_2(M/2) = 1 \text{ bits}. \quad (26)$$

Therefore, according to this coarse model, $\mathbf{h2}$ is expected to carry about 1 bit of information, for it halves the search space, when $\mathbf{h1}$ is not provided and the word embedding is sparse. However, the effect of $\mathbf{h2}$ after $\mathbf{h1}$ is less evident. In practice, results such as $\widehat{\Delta I}_2 = 0.7$, for the subset *gram1-adjective-to-adverb*, in Table I, seems to suggest that the effective search space is reduced to more than half that spotted by $\mathbf{h1}$, whereas $\widehat{\Delta I}_2 = 5.1$, for the subset *gram4-superlative*, suggests a much stronger reduction of this search space to less than 3% of the words spotted by $\mathbf{h1}$. Besides, from the same perspective, the negative result in Table II, $\widehat{\Delta I}_2 = -1.0$, for the *Lexicographic semant.* subset suggests that hint $\mathbf{h2}$ is misleading the guesser, doubling in size the search space already spotted by $\mathbf{h1}$. A simple illustration for how is it possible is presented in the Appendix.

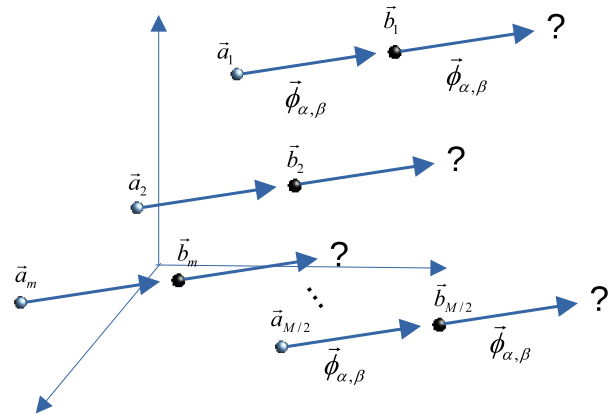


Fig. 5. Illustration of how start-words and end-words yield qualitatively different predictions when added to vector $\vec{\phi}_{\alpha, \beta}$. Start-words, represented by \vec{a} vectors, yield prediction close to \vec{b} vectors, representing end-words, whereas end-words yield predictions most probably far from other vectors. As a result, when guessing candidates are sorted according to their respective distances to respective predictions, the target \vec{b}_t is expected to be found among the first $M/2$ near neighbours.

VI. CONCLUSION

An approach for measuring the information content of hints in analogy tests was proposed. For conciseness, all experiments were done with a pre-trained set of GloVe 300-D vectors, whose performance in analogy tests is considered representative of most state-of-the-art embeddings. The test sets were either the Google Analogy Test Set, or the Bigger Analogy Test Set. All experiments corroborate the general perception noticed in publications since 2016 [10], [11], [12], [13], [10], that the analogy hint is much less relevant to the test performances than word proximity, even if word vectors do capture linguistic regularities. This was confirmed in all but one subset of tests, where the analogy hint had a negative effect, suggesting that sometimes analogy hints may even play the role of a disturbing noise, or a false hint for the guessing game. Indeed, the negative value of $\widehat{\Delta I}_2$, in Table II, for the last subset of tests, means that analogy hints hindered the target word search.

Results, when regarded solely in terms of accuracy, may suggest that the analogy hint is strongly relevant, at least for the purpose of accuracy gain in tests. However true, this is a fallacious result inasmuch as accuracy falsely amplifies the actual importance of **h2**. For instance, a result such as $\widehat{\Delta I}_1 = 16.8$, for the subset *gram6-nationality-adjective*, in Table I, indicates that **h1** reduced the search space from $M = 400,000$ to a set of spotted words with effective cardinality of about 4 words, yielding an accuracy of about 29.5%. After **h1**, **h2** has a very modest effect of further reducing the search space to a set whose effective cardinality is about 2 words, but the accuracy impressively increases to 68.8%. Indeed, the accuracy gain is amplified by the fact that it is inversely proportional to the effective size of the search space, already mostly reduced by the effect of **h1**.

In [15], another measure for the presence of linguistic relations in word embeddings was used, the Pairing Consistency Score (PCS), which quantifies the degree to which offset vectors $\vec{\phi}_{a,b}$ and $\vec{\phi}_{\alpha,\beta}$ are parallel above chance. Deviations between angle distributions were measured with Area Under the Curve (AUC) values. That is to say that AUC values in the range from 0.5 to 1.0 were used to evaluate the relevance of analogies in each set of experiments, and PCS values greater than 0.5 were taken as indicators that $\vec{\phi}_{a,b}$ and $\vec{\phi}_{\alpha,\beta}$ were parallel above chance. From the information content perspective, within-category shuffling does break analogies, so that the corresponding pseudo hint **h2** should carry null information content, which indeed corresponds to what suited experiments in [15]. Unfortunately, the resulting deviations of PCS from 0.5, which are proposed there as measures of linguistic relations relevance, are not easily translatable into accuracy gains due to these same linguistic relations, otherwise it would be an interesting matter for direct comparison with the approach proposed here.

As for indirect comparisons, however, from results presented in [15], the pairing consistency score is more deviated from 0.5 for the *Inflectional* subset of BATS, whereas $\widehat{\Delta I}_2 = 1.9$, in Table II, is the second highest estimated value of ΔI_2 for BATS, thus corroborating the relevance of

analogy in that subset. By contrast, PCS is larger for subset *Derivational* than for the *Encyclopedic* one, while the opposite is observed in terms of $\widehat{\Delta I}_2$, in Table II. Moreover, although PCS is close to 0.5 for tests with the *Lexicographic* subset, it is still greater than 0.5, which means that offsets $\vec{\phi}_{a,b}$ and $\vec{\phi}_{\alpha,\beta}$ are parallel above chance, on average, whereas in Table II we find a negative variation of information content due to hint **h2**.

These results are not inconsistent. In particular, for the *Lexicographic* subset, PCS value suggests a non-random parallelism between $\vec{\phi}_{a,b}$ and $\vec{\phi}_{\alpha,\beta}$, in that subset, whereas the negative $\widehat{\Delta I}_2$ suggests that $\vec{\phi}_{\alpha,\beta}$ is moving the search toward a neighbourhood with more target candidates than the prediction neighbourhood after **h1** alone. Indeed, because $\widehat{\Delta I}_2 = -1$ bit, we may even infer that $\vec{p}_2 = \vec{a} + \vec{\phi}_{\alpha,\beta}$ spots a portion of the search space with effective cardinality twice as larger as that spotted by $\vec{p}_1 = \vec{a}$, on average. Therefore, the two results do not contradict each other, for although $\widehat{\Delta I}_2$ also depends on the parallelism between $\vec{\phi}_{a,b}$ and $\vec{\phi}_{\alpha,\beta}$, as PCS, it also depends on the effective cardinality of spotted portions of the search space.

Therefore, the proposed approach allows a new complementary way of measuring the actual relevance of analogy hints, as compared to proximity hints. Besides, the vector offset method is just one kind of signal prediction model that can be used in word embeddings, and the proposed approach for information content estimation has the potential to be a useful tool in further studies regarding signal prediction in embeddings.

REFERENCES

- [1] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, and J. Dean. "Distributed representations of words and phrases and their compositionality." *Advances in neural information processing systems*, vol. 26, pp. 3111–3119, 2013, doi: 10.5555/2999792.2999959.
- [2] H. K. Kim, H. Kim, and S. Cho. "Bag-of-concepts: Comprehending document representation through clustering words in distributed representation." *Neurocomputing*, vol. 266, no. 29, pp. 336–352, 2017, doi: 10.1016/j.neucom.2017.05.046.
- [3] Z. Cui, X. Shi, and Y. Chen. "Sentiment analysis via integrating distributed representations of variable-length word sequence." *Neurocomputing*, vol. 187, no. 26, pp. 126–132, 2016, doi: 10.1016/j.neucom.2015.07.129.
- [4] S. Lim, H. Prade, and G. Richard. "Classifying and completing word analogies by machine learning." *International Journal of Approximate Reasoning*, vol. 132, pp. 1–25, 2021, doi: 10.1016/j.ijar.2021.02.002.
- [5] J. Choi and S.-W. Lee. "Improving fasttext with inverse document frequency of subwords." *Pattern Recognition Letters*, vol. 133, pp. 165–172, 2020, doi: 10.1016/j.patrec.2020.03.003.
- [6] T. Sousa, H. Gonalo Oliveira, and A. Alves. "Exploring different methods for solving analogies with portuguese word embeddings." In *9th Symposium on Languages, Applications and Technologies (SLATE 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020, doi: r 10.4230/OASfcs.SLATE.2020.9.
- [7] R. A. Stein, P. A. Jaques, and J. F. Valiati. "An analysis of hierarchical text classification using word embeddings." *Information Sciences*, vol. 471, pp. 216–232, 2019, doi: /10.1016/j.ins.2018.09.001.
- [8] X. Cao, J. Li, R. Wang, Y. W., Q. Niu, and J. Shi. "Calibrating glove model on the principle of zipf's law." *Pattern Recognition Letters*, vol. 125, pp. 715–720, 2019, doi: 10.1016/j.patrec.2019.07.015.
- [9] T. Mikolov, K. Chen, G. Corrado, and J. Dean. "Efficient estimation of word representations in vector space." *arXiv preprint arXiv:1301.3781*, 2013, doi: 10.48550/arXiv.1301.3781.
- [10] T. Linzen. "Issues in evaluating semantic spaces using word analogies." *arXiv preprint arXiv:1606.07736*, 2016, doi: 10.48550/arXiv.1606.07736.

[11] M. Faruqui, Y. Tsvetkov, P. Rastogi, and C. Dyer. “Problems with evaluation of word embeddings using word similarity tasks.” *arXiv preprint arXiv:1605.02276*, 2016, doi: 10.48550/arXiv.1605.02276.

[12] Y. Goldberg and O. Levy. “word2vec explained: deriving mikolov et al.’s negative-sampling word-embedding method.” *arXiv preprint arXiv:1402.3722*, 2014, doi: 10.48550/arXiv.1402.3722.

[13] O. Levy and Y. Goldberg. “Linguistic regularities in sparse and explicit word representations.” In *Proceedings of the eighteenth conference on computational natural language learning*, pp. 171–180, 2014, doi: 10.3115/v1/W14-1618.

[14] K. Lund and C. Burgess. “Producing high-dimensional semantic spaces from lexical co-occurrence.” *Behavior research methods, instruments, & computers*, vol. 28, no. 2, pp. 203–208, 1996, doi: 10.3758/BF03204766.

[15] L. Fournier, E. Dupoux, and E. Dunbar. “Analogies minus analogy test: measuring regularities in word embeddings.” *arXiv preprint arXiv:2010.03446*, 2020, doi: 10.48550/arXiv.2010.03446.

[16] D. JC MacKay. *Information theory, inference and learning algorithms*. Cambridge university press, 2003, doi: 10.5555/971143.

[17] C. E. Shannon. “A mathematical theory of communication.” *The Bell system technical journal*, vol. 27, no. 3, pp. 379–423, 1948, doi: 10.1002/j.1538-7305.1948.tb01338.x.

[18] J. Montalvão, J. Canuto, and E. Carvalho. “On the minimum probability of classification error through effective cardinality comparison.” *Journal of Communication and Information Systems*, vol. 31, no. 1, pp. 216–232, 2016, doi: 10.14209/jcis.2016.26.

[19] G. Miller. “Note on the bias of information estimates.” *Information theory in psychology: Problems and methods*, 1955.

[20] J. Pennington, R. Socher, and C. D. Manning. “Glove: Global vectors for word representation.” In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014, doi: 10.3115/v1/D14-1162.

[21] 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 *Proceedings of the NAACL Student Research Workshop*, pp. 8–15, 2016, doi: 10.18653/v1/N16-2002.

APPENDIX

A BRIEF ILLUSTRATION OF EFFECTIVE CARDINALITY AND INFORMATION CONTENT

Two main concepts used in this work are *effective cardinality* and *information content*. To illustrate these concepts, we consider a game where a (loaded) six-sided dice is thrown and a random variable X is associated to the corresponding outcomes, where $\Pr(X = 1) = 1/6$, $\Pr(X = 2) = 1/24$, $\Pr(X = 3) = 1/3$, $\Pr(X = 4) = 1/12$, $\Pr(X = 5) = 1/3$ and $\Pr(X = 6) = 1/24$.

The Shannon Information Content of an outcome x is defined in [16] in bits as

$$I(x) = \log_2 \frac{1}{\Pr(X = x)}$$

which turns out to be, in Shannon’s original work ([17], Section 7), the proposed *measure of how much “choice” is involved in the selection of the event or of how uncertain we are of the outcome*. The Information Content in this work is a short for the Shannon Information Content, which is also defined for a random categorical outcome, such as the choice of a word.

Before any hint is provided, a guesser has no reason to prefer any outcome, then he randomly chooses one out of the six faces at random. Therefore, the chances of each face being chosen by the guesser is $1/6$, and whatever his choice may be, the average accuracy of this guesser can be easily computed as being $1/6$, in spite of the non-uniform probability distribution of the loaded dice. Therefore, $I_0 = \log_2(6)$ bits quantifies the missing information for the guesser, even though

the probability distribution of the dice is not uniform. It comes from the fact that the unbalanced distribution is not known by the guesser.

Then, a first hint, **h1**, is provided, according to which outcomes from the loaded dice have an expected value around 3 and 4, which is true, for $E\{X\} = 3.5$. Suppose the guesser reacts to **h1** by updating his model of the dice distribution to $\Pr(X = 1) \approx 1/8$, $\Pr(X = 2) \approx 1/8$, $\Pr(X = 3) \approx 1/4$, $\Pr(X = 4) \approx 1/4$, $\Pr(X = 5) \approx 1/8$ and $\Pr(X = 6) \approx 1/8$. Indeed, this distribution roughly reflects what the guesser just learned from **h1**, and now, to maximize his chances of a hit, he chooses either face 3 or face 4. Therefore, the average accuracy of this guesser increases to $5/24$, then $2^{-I_1} = 5/24$, or $I_1 = \log_2(24/5) \approx 2.26$ bits, and $\Delta I_1 \approx 0.32$ bits.

At this point, the concept of *effective cardinality* can be conveniently introduced as follows: from the viewpoint of the guesser, in spite of the actual probability distribution of the loaded dice, the first problem was that of finding a single right answer in a set of $C_0 = 2^{I_0} = 6$ equally probable ones. Clearly enough, C_0 is the actual cardinality of this set. After **h1** was provided, this problem was modified to finding a single right answer in a (chimerical) set of $C_1 = 2^{I_1} = 4.8$ equally probable ones. In this case, C_1 is not an actual set cardinality, but it plays the role of an *effective cardinality* that intuitively reflects the difficulty level of the guessing game, in this context. More generally, effective cardinality can be a useful dual for the entropy concept itself, as explained in [18] and references therein.

Finally, to illustrate why ΔI_2 can be negative, consider that a second hint, **h2**, was provided, according to which outcomes from the loaded dice are more likely to be 4, 5 or 6, which is false! Now, not knowing that the hint is wrong, the guesser may update his former model to also include **h2**, as $\Pr(X = 1) \approx 1/13$, $\Pr(X = 2) \approx 1/13$, $\Pr(X = 3) \approx 3/13$, $\Pr(X = 4) \approx 4/13$, $\Pr(X = 5) \approx 2/13$ and $\Pr(X = 6) \approx 2/13$.

From the guesser perspective, To maximize his chances of a hit, in accordance to this updated model, his guess should be always face 4. However, because **h2** is a false hint, the guesser accuracy in lessened to $1/12$, or $I_2 = \log_2(12) \approx 3.58$ bits, and $\Delta I_2 \approx -1.32$ bits. Alternatively, one may consider that the guessing game was hardened by the misleading hint **h2**, from searching the right answer in a set of $C_1 = 4.8$ equally probable ones, to searching in a bigger virtual set of $C_2 = 12$ also equally likely answers. It is noteworthy that it is even harder than guessing the outcome of a fair dice with 6 faces, thus with cardinality C_0 . In other words, the original effective cardinality was expanded as a consequence of a false hint.