

Can Language Models Encode Perceptual Structure Without Grounding? A Case Study in Color

Mostafa Abdou*
University of Copenhagen

Artur Kulmizev
Uppsala University

Daniel Hershcovich
University of Copenhagen

Stella Frank
University of Trento

Ellie Pavlick
Brown University

Anders Søgaard
University of Copenhagen

Abstract

Pretrained language models have been shown to encode relational information, such as the relations between entities or concepts in knowledge-bases — (Paris, Capital, France). However, simple relations of this type can often be recovered heuristically and the extent to which models implicitly reflect topological structure that is grounded in world, such as perceptual structure, is unknown. To explore this question, we conduct a thorough case study on color. Namely, we employ a dataset of monolexic color terms and color chips represented in CIELAB, a color space with a perceptually meaningful distance metric.

Using two methods of evaluating the structural alignment of colors in this space with text-derived color term representations, we find significant correspondence. Analyzing the differences in alignment across the color spectrum, we find that warmer colors are, on average, better aligned to the perceptual color space than cooler ones, suggesting an intriguing connection to findings from recent work on efficient communication in color naming. Further analysis suggests that differences in alignment are, in part, mediated by collocationality and differences in syntactic usage, posing questions as to the relationship between color perception and usage and context.

1 Introduction

Without grounding or interaction with the world, language models (LMs) learn representations that encode various aspects of formal linguistic structure (e.g., morphosyntax (Tenney et al., 2019)) and semantic information (e.g., lexical similarity (Reif et al., 2019a)). Beyond this, it has been suggested that text-only training data is enough for LMs to also acquire factual and relational information about the world (Davison et al., 2019; Petroni et al., 2019). This includes, for instance, some




Figure 1: Right: Color orientation in 3d CIELAB space. Left: linear mapping from BERT (CC, see §2) color term embeddings to the CIELAB space.

features of concrete and abstract concepts, such as objects’ attributes and affordances (Forbes et al., 2019b; Weir et al., 2020). Furthermore, the representational geometry of LMs has been found to naturally reflect human lexical similarity and relatedness judgements, as well as analogy relationships (Chronis and Erk, 2020). However, the extent to which these models reflect the structures that exist in humans’ perceptual world—such as the topology of visual perception (Chen, 1982), the structure of the color spectrum (Ennis and Zaidi, 2019; Provenzi, 2020), or of odour spaces (Rossiter, 1996; Chastrette, 1997)—is not well-understood.

If LMs are indeed able to capture such topologies—in some domains, at least—it would mean that these structures are a) somehow reflected in language and, thereby, encoded in the textual training data on which models are trained, and b) learnable using models’ current training objectives and architectural inductive biases. To the extent they are not, the question becomes whether the information is not there in the data, or whether model and training objective limitations are to blame. Certainly, this latter point relates to an ongoing debate regarding what exactly language models can be expected to learn from ungrounded form alone (Bender and Koller, 2020; Bisk et al., 2020; Merrill et al., 2021). While there have been many inter-

For correspondence: {abdou, soegaard}@di.ku.dk

esting theoretical debates around this topic, few studies have tried to address this question empirically.

In this paper, we conduct a case study on color. Indeed, color perception in humans and its relation to speakers’ use of color terms has long been the subject of studies in cognitive science (Kay and McDaniel, 1978; Berlin and Kay, 1991; Regier et al., 2007; Kay et al., 2009). To this end, spaces have been defined in which Euclidean distances between related colors are correlated with reported perceptual differences.¹ In addition, the semantics of color terms have long been understood to hold particular linguistic significance, as they are theorised to be subject to universal constraints that arise directly from the neurophysiological mechanisms and properties underlying visual perception and cognition (Kay and McDaniel, 1978; Berlin and Kay, 1991; Kay et al., 1991).² Due to these factors, color offers a useful test-bed for investigating whether or not structural information about the topology of the perceptual world might be encoded in linguistic representations.

To explore this in detail, we employ a dataset of English color terms and their corresponding color chips³, the latter of which are represented in CIELAB — a perceptually uniform color space. In addition to the color chip CIELAB coordinates, we extract linguistic representations for the corresponding color terms. With these two representations in mind (see Figure 1 for a demonstrative plot from our experiments), we employ two methods of measuring structural correspondence, with which we evaluate the alignment between the two spaces. Figure 2 shows an illustration of the experimental setup. We find that the structures of various language model representations show alignment with the structure of the CIELAB space, demonstrating that some approximation of perceptual color space topology can indeed be learned from text alone.

¹The differences between color stimuli which are perceived by human observers.

²These theories have been contested by work arguing for linguistic relativism (cf. the *Sapir-Whorf Hypothesis*), which emphasizes the arbitrariness of language and the relativity of semantic structures and minimizes the role of universals. Such critiques have, however, been accommodated for in the Berlin & Kay paradigm (Berlin and Kay, 1991), the basic assumptions of which, such as the existence of at least some perceptually-determined universal constraints on color naming, remain widely accepted.

³Each chip is a unique color sample from the Munsell chart, which is made up of 330 such samples which cover the space of colors perceived by humans. See §2.

We also show that part of this distributional signal is learnable by simple models — e.g. models based on pointwise mutual information (PMI) statistics — although large-scale language model pretraining (e.g., BERT) encodes the topology markedly better.

Analysis shows that larger language models align better than smaller ones and that much of the variance in CIELAB space can be explained by low-dimensional subspaces of LM-induced color term representations. To better understand the results, we also analyse the differences in alignment across the color spectrum, observing that warm colors are generally better aligned than cool ones. Further investigation reveals a connection to findings reported in work on communication efficiency in color naming, which posits that warmer colors are communicated more efficiently. Finally, we investigate various corpus statistics which could influence alignment, finding that a measure of color term collocationality based on PMI statistics corresponds to lower alignment, while the entropy of a color term’s dependency relation distribution (i.e. terms occurring as adjectival modifiers, nominal subjects, etc.) and how often it occurs as an adjectival modifier correspond to a stronger one.

2 Methodology

Color data We employ the Color Lexicon of American English, which provides extensive data on color naming. The lexicon consists of 51 monolexemic color name judgements for each of the 330 Munsell Chart color chips⁴ (Lindsey and Brown, 2014). The color terms are solicited through a free-naming task, resulting in 122 terms.

Perceptual color space Following previous work (Regier et al., 2007; Zaslavsky et al., 2018; Chaabouni et al., 2021), we map colors to their corresponding points in the 3D CIELAB space, where the first dimension L expresses lightness, the second A expresses position between red and green, and the third B expresses the position between blue and yellow. Distances between colors in the space correspond to their perceptual difference.

Language models Our analysis is conducted on three widely used language models (LMs): BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019), both of which employ a masked language modelling objective, and ELECTRA (Clark et al.,

⁴<http://www1.icsi.berkeley.edu/wcs/images/jrus-20100531/wcs-chart-4x.png>

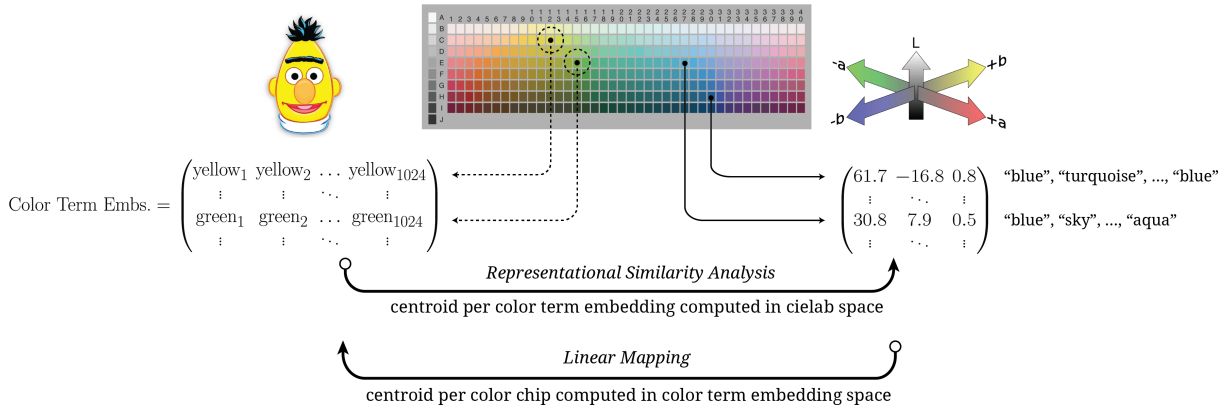


Figure 2: Our experimental setup. In the center is a Munsell color chart. Each chip in the chart is represented in the CIELAB space (right) and has 51 color term annotations. Color term embeddings are extracted through various methods. In the Representation Similarity Analysis experiments, a corresponding color chip centroid is computed in the CIELAB space. In the Linear Mapping experiments, a color term embedding centroid is computed per chip.

2020), which is trained instead with a discriminative token replacement detection objective.⁵

Baselines In addition to the aforementioned language models, we consider two different baselines:

- PMI statistics, which are computed⁶ for the color terms in common crawl, using window sizes of 1 (pmi-1), 2 (pmi-2), and 3 (pmi-3). The result is a vocabulary length vector quantifying the likelihood of co-occurrence of the color term with every other vocabulary item in within that window.
- Word-type FastText embeddings trained on Common Crawl (Bojanowski et al., 2017).

Representation Extraction We follow Bommasani et al. (2020) and Vulić et al. (2020) in defining configurations for the extraction of word-type representations from LM hidden states. In the first configuration (NC), a color term is encoded without context, with the appropriate delimiter tokens attached (e.g. [CLS] red [SEP] for BERT). In the second, S sentential contexts that include the color term are encoded and the hidden states representing these contexts are mean pooled. These S contexts are either randomly sampled from common crawl (RC), or deterministically generated to allow for control over contextual variation (CC). If a color term is split by an LM’s tokenizer into more than one token, subword token encodings are averaged over. For each color term and configuration,

⁵bert-large-uncased; roberta-large; electra-large-discriminator

⁶Using Hyperwords: <https://bitbucket.org/omerlevy/hyperwords>

an embedding vector of hidden state dimension d_{LM} is extracted per layer, per model.

Controlled context To control for the effect of variation in the sentence contexts used to construct color term representations, we employ a templative approach to generate a set of identical contexts for all color terms. When generating controlled contexts, we create three frames in which the terms can appear:

- COPULA: the <obj> is <col>
- POSSESSION: i have a <col> <obj>
- SPATIAL: the <col> <obj> is there

We use these frames in order to limit the contextual variation across colors (<col>) and to isolate their representations amidst as little semantic interference as possible, all while retaining a naturalistic quality to the input. We also aggregate over numerous object nouns (<obj>), which the color terms are used to describe. We select objects from the McRae et al. (2005) data which are labelled in the latter as plausibly occurring in many colors and which are stratified across 13 category sets, e.g. *fan* ∈ APPLIANCES, *skirt* ∈ CLOTHING, etc. Collapsing over categories, we generate sentences combinatorially across frames, objects and color terms, resulting in $3 \times 122 \times 18 = 6588$ sentences, 366 per term.

3 Evaluation

We employ two complimentary evaluation methods to gauge the correspondence of the color term text-derived representations to the perceptual color

space. The first, Representation Similarity Analysis (RSA), is non-parametric and uses pairwise comparisons of stimuli to provide a measure of the global topological alignment between two spaces. The second employs a learned linear mapping, evaluating the extent to which two spaces can be aligned via transformation (rotation, scaling, etc.).

RSA (Kriegeskorte et al., 2008) is a method of relating different representational modalities, which was first employed in neuroscientific studies. RSA abstracts away from activity patterns themselves (e.g. neuron values in representational vectors) and instead computes representational (dis)-similarity matrices (RSMs), which characterize the information carried by a given representation method through global (dis)-similarity structure. Kendall’s rank correlation coefficient (τ) is computed between RSMs derived from the two spaces, providing a summary statistic indicative of the overall representational alignment between them. RSA is non-parametric and therefore circumvents many of the various methodological weaknesses associated with the probing paradigm (Belinkov, 2021).

For each color term, we compute a centroid in the CIELAB space following the approach described in Lindsey and Brown (2014). Each centroid is defined as the average CIELAB coordinate of the samples (i.e. color chips) that were named with the corresponding term (across the 51 subjects). This results in N parallel points in the color term embedding and perceptual color spaces, where N is the number of color terms considered. For our analysis, we exclude color terms used less frequently than a cutoff $f = 100$ in the color lexicon, leaving us with the 18 most commonly used color terms.⁷ We then separately construct an $N \times N$ RSM for each of the LM spaces and for CIELAB. Each cell in the RSM corresponds to the similarity between the activity patterns associated with pairs of experimental conditions $n_i, n_j \in N$.

For the color term embedding space, we employ Pearson’s correlation coefficient (r) as a similarity measure between each pair of embeddings $n_i, n_j \in N$. For the CIELAB space, we elect to use the following method, per Regier et al.’s (2007) suggestion: $sim(n_i, n_j) = \exp(-c \times [dist(n_i, n_j)]^2)$, where c is a scaling factor (set to 0.001 in all ex-

⁷This includes all color terms which are considered "basic" (red, blue, etc.), and commonly used "derived" terms (pink, gray, turquoise, maroon, etc.), but excludes the rest which are only infrequently used as color terms (forest, puke, dew, seafoam, etc.). See appendix A for full list of colors included.

periments reported here) and $dist(n_i, n_j)$ is the CIELAB distance (ΔE_{CMC}^*)⁸ between chips n_i and n_j . This similarity measure is derived from the psychological literature on categorization and is meant to model the assumption that beyond a certain distance colors appear entirely different, so that increasing the distance has no further effect on dissimilarity. Finally, we report the mean Kendall’s τ between the color term embedding and color space RSMs. We also report τ per color term (i.e. per row in the RSM), which corresponds to how well-aligned each individual color term is.

Linear mapping We train regularised linear regression models to map from color term embedding space $X \in \mathbb{R}^{n \times d_{LM}}$ to CIELAB space $Y \in \mathbb{R}^{n \times 3}$, minimising $\mathcal{L}(W; \alpha) = \|XW - Y\|_2^2 + \alpha \|W\|_1$, where $W \in \mathbb{R}^{3 \times d_{LM}}$ is a linear map and α is the lasso regularization hyper parameter. We vary α across a wide range of settings to examine the effect of probe complexity, which we measure using the nuclear norm of the linear projection matrix $W \in \mathbb{R}^{\phi \times \iota}$; $\|W\|_* = \sum_{i=1}^{\min(\phi, \iota)} \sigma_i(W)$, where $\sigma_i(W)$ is the i th singular value of W (Pimentel et al., 2020). The fitness of the regressors, evaluated using n -fold cross-validation ($n = 6$) indicates the alignability of the two spaces, given a linear transformation. Centroids corresponding to each Munsell color chip are computed in the color term embedding space via the weighted mean of the embeddings of the 51 terms used to label it. As in the RSA experiments, terms occurring less frequently than the cutoff ($f = 100$) are excluded. For evaluation, we compute the average (across splits and datapoints) proportion of explained variance as well as the ranking of a predicted color term embedding according to the Pearson distance ($1 - r$) to gold.

Control task As proposed by Hewitt and Liang (2019), we construct a random control task for the linear mapping experiments, wherein we randomly swap each color chip’s CIELAB code for another. This is meant to break the mapping between the color chips and their corresponding terms. Control task results are reported as the mean of 10 different random re-mappings. We report probe *selectivity*, which is defined as the difference between proportion of explained variance in the standard experimental condition and in the control task (He-

⁸We use the `colormath` Python package, setting illuminant to C, and assuming 2 degree standard observer.

Model	NC				RC				CC			
	RSA		lin. map		RSA		lin. map		RSA		lin. map	
	max	mean	max	mean	max	mean	max	mean	max	mean	max	mean
BERT	0.16*	0.01 \pm 0.09	0.75	0.73 \pm 0.01	0.26 \dagger	0.20 \pm 0.03	0.74	0.73 \pm 0.08	0.24 \dagger	0.19 \pm 0.03	0.76	0.75 \pm 0.05
RoBERTa	0.33 \S	0.02 \pm 0.11	0.75	0.73 \pm 0.01	0.20*	0.14 \pm 0.04	0.74	0.73 \pm 0.01	0.19*	0.14 \pm 0.04	0.77	0.76 \pm 0.09
ELECTRA	0.13	0.01 \pm 0.08	0.75	0.64 \pm 0.13	0.25 \dagger	0.19 \pm 0.05	0.75	0.73 \pm 0.01	0.23 \dagger	0.16 \pm 0.04	0.78	0.76 \pm 0.01

Table 1: Results for the RSA experiments show max and mean (across layers) Kendall’s τ ; correlations that are significantly non-zero are marked with *, \dagger and \S for $p < 0.05$, < 0.01 and < 0.001 respectively. Results for the linear mapping experiments show max and mean selectivity.

witt and Liang, 2019). We run similar control for the RSA experiments, where the CIELAB space centroids are randomly shuffled.

4 Results

Table 1 shows the max, mean, and standard deviation (across layers) of alignment scores for each of the LMs, per alignment method and setting. For RSA, we observe significant correlations across all configurations: most LM layers show a topological alignment with color space. Notably, this is also true for the static embeddings and for one of the PMI baselines (Table 2). Although some variance is observed,⁹ the presence of significant correlations is telling, given the small sample size (18). Furthermore, randomly permuting the color space centroids leads to RSA correlations that are non-significant for all setups ($p > 0.05$), which lends further credence to models’ alignment with CIELAB structure.

Figure 3 shows the breakdown of correlations per color term for the three LMs under CC, as well as for fastText. We find that this ranking of color terms is largely stable across models and layer. Full RSMs for all models and CIELAB are in appendix C. The RSMs show evidence of the higher correlations for colors like violet, orange, and purple, being driven by general clusterings of similarity/dissimilarity. For instance, for both the CIELAB and CC BERT RSMs, violet’s top *nearest* neighbors include purple, lavender, pink, and orange, and its *furthest* neighbors include aqua, olive, black, and gray. Correlations do not, however, appear to be driven by consistently aligned partial orderings within the clusters. In addition, we compute RSA correlations between the different

⁹In particular, results for NC show large variances across layers. The mean correlation across layers in this setup is near zero, even though max correlations for BERT and RoBERTa are significant; this is unsurprising, however, as the LM has likely never encountered single color term tokens in isolation (cf. Bommasani et al. (2020))

Model	RSA	lin. map
pmi-1	0.14	0.72
pmi-2	0.11	0.70
pmi-3	0.17*	0.71
fastText	0.23*	0.72

Table 2: Baseline results. RSA results show Kendall’s τ ; results with * are significantly non-zero ($p < 0.05$). Linear mapping results show selectivity.

models. Results show that NC embeddings have low alignment to all others (details in appendix B).

For the linear mapping experiments, we observe the highest selectivity scores for CC (Table 1, right) compared to NC and RC (Table 1, left, middle) and baselines (Table 2). This validates our intuition that controlling for variation in sentence context would reveal increased alignment to color space.

Furthermore, we observe that, over the full range of probe complexities for the experimental condition and the control task (described as in §3), all models demonstrate high selectivity (see G for full results). It is, therefore, safe to attribute the fitness of the probes to information encoded in the color term representations, rather than to memorization. In terms of individual colors, Figure 4a depicts the ranking of predicted CIELAB codes per Munsell color chip for BERT (CC). We find that these results are largely stable across models and layers (see appendix F for full set of results and for reference chart). Also, we observe that clusterings of chips with certain modal color terms (*green*, *blue*) show worse rankings than the rest.

5 Analysis and Discussion

Having demonstrated the existence of models’ alignment to CIELAB across various configurations, we now present an analysis and discussion of these results.

Dimensionality of color subspace Previous work has shown that linguistic information such as




Figure 3: RSA results (Kendal’s τ) broken down by color term for each of the LMs under the CC configuration and for the fastText baseline.

part-of-speech category, dependency relation type, and word sense, is expressed in low-dimensional subspaces of language model representations (Reif et al., 2019b; Durrani et al., 2020; Hernandez and Andreas, 2021). We investigate the dimensionality of the subspace required to predict the CIELAB chip codes from the term embeddings, following the methodology of Durrani et al. (2020). Averaging over the three predicted CIELAB dimensions, we rank the linear mapping coefficients (from the experiments described in §2), sorting the weights by their absolute values in descending order. Results (appendix H) show that across models and layers, ~ 0.4 of the variance in the CIELAB chip codes can be explained by assigning 95% of the weights to ~ 10 dimensions. 30–40 dimensions are sufficient to explain ~ 0.7 of the variance, nearly the proportion of variance explained by the full representations (Table 1).

Model	RSA max	RSA mean	lin. map. max	lin. map. mean
BERT-mini	0.077	0.043 \pm 0.340	0.729	0.582 \pm 0.291
BERT-small	0.106	0.070 \pm 0.191	0.734	0.598 \pm 0.294
BERT-medium	0.097	0.057 \pm 0.035	0.739	0.654 \pm 0.221
BERT-base	0.162*	0.092 \pm 0.058	0.740	0.677 \pm 0.182

Table 3: Results for the four smaller BERT models. RSA results (left) show max and mean (across layers) Kendall’s correlation coefficient (τ). Correlations that are significantly non-zero are indicated with: * : $p < 0.05$. Results for the Linear Mapping experiments (right) show max and mean selectivity. Standard deviation across layers is included with the mean results.

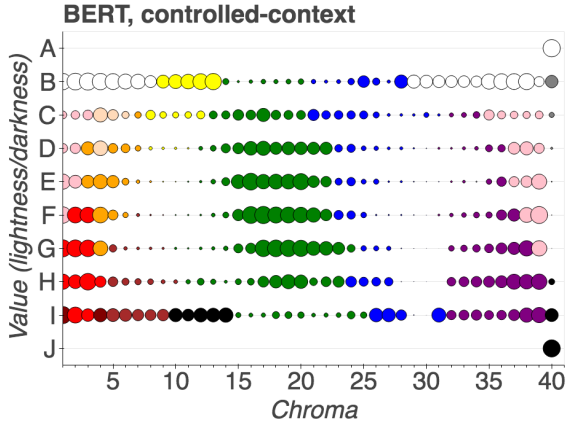
Effect of model size We also evaluate the effect of model size on alignment by testing four smaller BERT (CC) models¹⁰ using the same setup described above. The results (table 3) show that alignment as measured by both RSA and linear mapping progressively increases with model size,

¹⁰for details see appendix I


meaning that that with growing complexity, model representational geometry of color terms moves towards isomorphism to CIELAB.

Color temperature In Figures 3 & 4a we observe that on average, warmer colors (*yellow, orange, red, etc.*) show a closer alignment than cooler ones (*blue, green, etc.*). In recent work, Gibson et al. (2017) reported that the former are on average communicated more efficiently (see next paragraph) than the latter, across languages. This is attributed to warmer colors being more prevalent as colors of behaviorally relevant items in the environment — salient objects — compared to cooler ones, which occur more often as background colors. To verify this observation, we partition the space of chips into two (see appendix D for details) and compute the average explained variance across warm and cool colors. The results (see appendix D for plots) show that, term embeddings of warm colors are better aligned to CIELAB than those of cool ones, across models and configurations. This is consistent with the bias described in Gibson et al. (2017), which we conjecture might be filtering through into the distributional statistics of (color terms in) textual corpora, influencing the representations learned by various methods which leverage these statistics.

Connection to listener surprisal Gibson et al. (2017)’s findings are based on the application of an information theoretic analysis to color naming, framing it as a communication game where a speaker has a particular color chip c in mind and uses a word w to indicate it then a listener has to correctly guess c , given w . Communication efficiency is measured through surprisal, S , which in this setting corresponds to the average number of guesses an optimal listener takes to arrive at the correct color chip. We calculate $S(c)$ for each



(a) Each circle on the chart represents the ranking of the predicted color chip when ranked according to Pearson distance from gold (larger circle \cong higher/better ranking).



(b) Each circle on the chart represents a color chip’s surprisal score (larger circle \cong higher score).

Figure 4: (a) shows linear mapping results for BERT, under the CC configuration, broken down by Munsell color chip; (b) shows surprisal per chip. Circle colors reflect the modal color term assigned to the chips.

chip in the color lexicon. Surprisal is defined as $S(c) = \sum_w P(w|c) \cdot \log\left(\frac{1}{P(c|w)}\right)$, where $P(w|c)$ is the probability that a color c gets labeled as w and $P(c|w)$ is computed using Bayes Theorem. Here, $P(w)$ represents how often a particular word gets used across the color space (and participants), and $P(c)$ is a uniform prior. Figure 4b shows surprisal per chip. High surprisal chips correspond to a lower color naming consensus among speakers, meaning that a more variable range of terms is used for these (color) contexts. We hypothesize that this could be reflected in the representations of color terms corresponding to high surprisal chips. To test this, we compute Spearman’s correlation (ρ) between a chip’s regression score (predicted color chip code ranking) and its surprisal. We find significant Spearman’s rank correlation between lower

ranking and higher surprisal for all LMs under all configurations ($0.12 \leq \rho \leq 0.17$, $p < 0.05$).

What factors predict color space alignment?

Given that LMs are trained exclusively on text corpora, we hypothesize that alignment between their embeddings and CIELAB is influenced by corpus usage statistics. To determine which factors could predict alignment score, we extract color term log frequency, part-of-speech tag (POS), dependency relation (DREL), and dependency tree head (HEAD) statistics for all color terms from a dependency-parsed (Straka et al., 2016) common crawl corpus. In addition to this, we compute, per color term, the entropy of its normalised PMI distribution (pmi-col , see §2) as a measure of collocation.¹¹ We then fit a Linear Mixed Effects Model (Gałeczki and Burzykowski, 2013) to the features listed above, with RSA score (Table 1) as the response variable, and model type as a random effect.

We follow a multi-level step-wise model building sequence, where a baseline model is first fit with color term log frequency as a single fixed effect. A model which includes pmi-col as an additional fixed effect is then fit, and these two terms are included as control predictors in all later models. Following this, we compute POS, DREL, and HEAD lemma distribution entropies per color term (pos-ent , dprel-ent , head-ent). Higher entropies indicate that the term is employed in more diverse contexts with respect to those categories. Following entropy computation, we separately fit models including each three entropy statistic features. Finally, we calculate the proportion of: POS tags that are adjectives, adj-prop ; DRELS that are adjectival modifiers, amod-prop ; and those that are copulas, cop-prop . The first two evaluate the effect of a color term occurring more or less often as an adjectival modifier, while the latter tests the hypothesis that assertions such as *The banana is yellow* could provide indirect grounding (Merrill et al., 2021), thereby leading to higher alignment. Including the entropy term which led to the best fit (dprel-ent) in the previous level, models are fit including terms for each of the proportion statistics. Model comparison is carried out by computing the log likelihood ratio between models that differ in a single term. See appendix J for model details.

¹¹Low entropy reflects frequent co-occurrence with a small subset of the vocabulary and high entropy the converse.

Results show that:

- `pmi-col` significantly improves fit above log frequency and has a negative coefficient, meaning that terms that occur in more fixed collocations are less aligned to the perceptual space. Intuitively, this makes sense as the color terms in many collocations such as e.g. *Red Army* or *Black Death* are employed in contexts which are largely metaphorical rather than attributive or descriptive.
- `deprel-ent` and `head-ent` (but not `pos-ent`) lead to a significantly improved fit compared to the control predictors; we observe positive coefficients for both, indicating RSA score is higher for terms that occur in more varied syntactic dependency relations and modify a more diverse set of syntactic heads. This suggests that occurring in a more diverse set of contexts might be beneficial for robust representation learning, in correspondence with the idea of sample diversity in the active learning literature (Brinker, 2003; Yang et al., 2015). `pos-ent`'s lack of significance, on the other hand, indicates that the degree of specification offered by the POS tagset might be too coarse to meaningfully differentiate between color terms, e.g. nouns can occur in a variety of DREs such as subjects, objects, oblique modifiers (per the Universal Dependencies (Nivre et al., 2020)).
- out of the proportion statistics, only the `amod-prop` term improves fit; it has a positive coefficient, thus color terms occurring more frequently as adjectival modifiers show higher scores. `adj-prop` is not significant, providing further evidence for the POS tagset's level of granularity being too coarse. Finally, as `cop-prop` is not significant, it appears that occurring more frequently in assertion-like copula constructions does not confer an advantage in terms of alignment to perceptual structure.

Vision-and-Language models In a preliminary set of experiments, we evaluated multi-modal Vision-and-Language models (VisualBERT (Li et al., 2019) and VideoBERT (Sun et al., 2019)), finding no major differences in results from the text-only models presented in this study.

6 Related Work

Distributional word representations have long been theorized to capture various types of information about the world (Schütze, 1992). Early work in this regard employed semantic similarity and relatedness datasets to measure alignment to human judgements (Agirre et al., 2009; Bruni et al., 2012; Hill et al., 2015). Rubinstein et al. (2015), however, question whether the distributional hypothesis is equally applicable to all types of semantic information, finding that taxonomic properties (such as animacy) are better modelled than attributive ones (color, size, etc.). To a similar end, Lucy and Gauthier (2017) analyze how well distributional representations encode various aspects of grounded meaning. They investigate whether language models would “*be worse off for not having physically bumped into walls before they hold discussions on wall-collisions?*”, finding that perceptual features are poorly modelled compared to encyclopedic and taxonomic ones.

More recently, several studies have asked related questions in the context of language models. For example, Davison et al. (2019) and Petroni et al. (2019) mine LMs for factual and commonsense knowledge by converting knowledge base triplets into cloze statements that are used to query the models. In a similar vein, Forbes et al. (2019a) investigate LM representations' encoding of object properties (e.g., *oranges are round*), and affordances (e.g. *oranges can be eaten*), as well as the interplay between the two. Weir et al. (2020) demonstrate that LMs can capture *stereotypic tacit assumptions* about generic concepts, showing that they are adept at retrieving concepts given their associated properties (e.g., **bear** given *A ___ has fur, is big, and has claws.*). Similar to other work, they find that LMs better model encyclopedic and functional properties than they do perceptual ones.

In an investigation of whether or not LMs are able to overcome reporting bias, Schwartz and Choi (2020) extract all sentences in Wikipedia where one of 11 color terms modifies a noun and test how well predicted the color term is when it is masked. They find that LMs are able to model this relationship between concepts and associated colors to a certain extent, but are prone to over-generalization. Finally, Ilharco et al. (2020) train a probe to map LM representations of textual captions to paired visual representations of image patches, in order to evaluate how useful the former are for discerning

between different visual representations. They find that many recent LMs yield representations that are effective at retrieving semantically-aligned image patches, but still far under-perform humans.

7 Outlook

It is commonly held that the learning of phenomena which rely on sensory perception is only possible through direct experience. Indeed, the view that people born blind could not be expected to acquire coherent knowledge about colors has been prevalent since at least the empiricist philosophers (Locke, 1847; Hume, 1938) and still holds currency (Jackson, 1982). Nevertheless, recent research highlighting the contribution of language and of semantic associations between concepts towards learning has demonstrated that the congenitally blind do in fact show a striking understanding of both color similarity (Saysani et al., 2018) and object colors (Kim et al., 2020).

This paper investigated whether representations of color terms that are derived from text only express a degree of isomorphism to the structure of humans’ perceptual color space.¹² Results from our experiments evidenced that such a topological correspondence exists. Notably, color term representations based on simple co-occurrence statistics already demonstrated correspondence; those extracted from language models aligned more closely. We observed that warm colors, on average, show more alignment than cooler ones, linking to recent findings on communication efficiency in color naming (Gibson et al., 2017).

Further analysis based on surprisal — an information theoretic measure, used to evaluate how efficiently a color is communicated between a speaker and a listener — revealed a correlation between lower topological alignment and higher color chip surprisal, suggesting that the kind of contexts a color occurs in play a role in determining alignment. Exploring this, we tested a set of color term corpus-derived statistics for how well they predict alignment, finding that a measure of a color term’s collocationality corresponds to lower alignment, while the entropy of its dependency relation distribution and it occurring more frequently as an adjectival modifier correspond to closer alignment.

¹²Clearly, complete isomorphism is rather unlikely: language in general, and color terms by extension, are far from being simply denotational, and language interacts with and is influenced by a myriad of factors besides perception.

Our results and analyses present empirical evidence of topological alignment between text-based color term representations and perceptual color spaces. With respect to the debate started by Bender and Koller (2020), we hope that this work offers a modest step towards furthering our understanding of the kinds of “meaning” we expect language models to acquire, with and without grounded or embodied learning approaches, and that it will provide motivation for further work in this direction.

Acknowledgements

We would like to thank Vinit Ravishankar and Mitja Nikolaus for their feedback and comments. Mostafa Abdou and Anders Søgaard are supported by a Google Focused Research Award and a Facebook Research Award.

References

- Eneko Agirre, Enrique Alfonseca, Keith Hall, Jana Kravalova, Marius Pasca, and Aitor Soroa. 2009. A study on similarity and relatedness using distributional and wordnet-based approaches.
- Yonatan Belinkov. 2021. Probing classifiers: Promises, shortcomings, and alternatives. *arXiv preprint arXiv:2102.12452*.
- Emily M. Bender and Alexander Koller. 2020. *Climbing towards NLU: On meaning, form, and understanding in the age of data*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 5185–5198, Online. Association for Computational Linguistics.
- Brent Berlin and Paul Kay. 1991. *Basic color terms: Their universality and evolution*. Univ of California Press.
- Yonatan Bisk, Ari Holtzman, Jesse Thomason, Jacob Andreas, Yoshua Bengio, Joyce Chai, Mirella Lapata, Angeliki Lazaridou, Jonathan May, Aleksandr Nisnevich, et al. 2020. Experience grounds language. *arXiv preprint arXiv:2004.10151*.
- Piotr Bojanowski, Edouard Grave, Armand Joulin, and Tomas Mikolov. 2017. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5:135–146.
- Rishi Bommasani, Kelly Davis, and Claire Cardie. 2020. *Interpreting Pretrained Contextualized Representations via Reductions to Static Embeddings*. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4758–4781, Online. Association for Computational Linguistics.

- Klaus Brinker. 2003. Incorporating diversity in active learning with support vector machines. In *Proceedings of the 20th international conference on machine learning (ICML-03)*, pages 59–66.
- Elia Bruni, Gemma Boleda, Marco Baroni, and Nam-Khanh Tran. 2012. Distributional semantics in technicolor. In *Proceedings of the 50th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 136–145.
- Rahma Chaabouni, Eugene Kharitonov, Emmanuel Dupoux, and Marco Baroni. 2021. Communicating artificial neural networks develop efficient color-naming systems. *Proceedings of the National Academy of Sciences*, 118(12).
- M Chastrette. 1997. Trends in structure-odor relationship. *SAR and QSAR in Environmental Research*, 6(3-4):215–254.
- Lin Chen. 1982. Topological structure in visual perception. *Science*, 218(4573):699–700.
- Gabriella Chronis and Katrin Erk. 2020. [When is a bishop not like a rook? when it’s like a rabbi! multi-prototype BERT embeddings for estimating semantic relationships](#). In *Proceedings of the 24th Conference on Computational Natural Language Learning*, pages 227–244, Online. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. 2020. Electra: Pre-training text encoders as discriminators rather than generators. *arXiv preprint arXiv:2003.10555*.
- Joe Davison, Joshua Feldman, and Alexander Rush. 2019. [Commonsense knowledge mining from pre-trained models](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 1173–1178, Hong Kong, China. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Nadir Durrani, Hassan Sajjad, Fahim Dalvi, and Yonatan Belinkov. 2020. [Analyzing individual neurons in pre-trained language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4865–4880, Online. Association for Computational Linguistics.
- Robert J Ennis and Qasim Zaidi. 2019. Geometrical structure of perceptual color space: mental representations and adaptation invariance. *Journal of vision*, 19(12):1–1.
- Maxwell Forbes, Ari Holtzman, and Yejin Choi. 2019a. Do neural language representations learn physical commonsense? *arXiv preprint arXiv:1908.02899*.
- Maxwell Forbes, Christine Kaeser-Chen, Piyush Sharma, and Serge Belongie. 2019b. [Neural naturalist: Generating fine-grained image comparisons](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 708–717, Hong Kong, China. Association for Computational Linguistics.
- Andrzej Gałeccki and Tomasz Burzykowski. 2013. Linear mixed-effects model. In *Linear Mixed-Effects Models Using R*, pages 245–273. Springer.
- Edward Gibson, Richard Futrell, Julian Jara-Ettinger, Kyle Mahowald, Leon Bergen, Sivalogeswaran Ratnasingam, Mitchell Gibson, Steven T Piantadosi, and Bevil R Conway. 2017. Color naming across languages reflects color use. *Proceedings of the National Academy of Sciences*, 114(40):10785–10790.
- Evan Hernandez and Jacob Andreas. 2021. The low-dimensional linear geometry of contextualized word representations. *arXiv preprint arXiv:2105.07109*.
- John Hewitt and Percy Liang. 2019. Designing and interpreting probes with control tasks. *arXiv preprint arXiv:1909.03368*.
- Felix Hill, Roi Reichart, and Anna Korhonen. 2015. Simlex-999: Evaluating semantic models with (genuine) similarity estimation. *Computational Linguistics*, 41(4):665–695.
- David Hume. 1938. *An Abstract of a Treatise of Human Nature, 1740*. CUP Archive.
- Gabriel Ilharco, Rowan Zellers, Ali Farhadi, and Hananeh Hajishirzi. 2020. Probing text models for common ground with visual representations. *arXiv preprint arXiv:2005.00619*.
- Frank Jackson. 1982. Epiphenomenal qualia. *The Philosophical Quarterly (1950-)*, 32(127):127–136.
- Paul Kay, Brent Berlin, Luisa Maffi, William R Merrifield, and Richard Cook. 2009. *The world color survey*. CSLI Publications Stanford, CA.
- Paul Kay, Brent Berlin, and William Merrifield. 1991. Biocultural implications of systems of color naming. *Journal of Linguistic Anthropology*, 1(1):12–25.
- Paul Kay and Chad K McDaniel. 1978. The linguistic significance of the meanings of basic color terms. *Language*, pages 610–646.

- Judy Sein Kim, Brianna Aheimer, Verónica Montané Manrara, and Marina Bedny. 2020. Shared understanding of color among congenitally blind and sighted adults.
- Nikolaus Kriegeskorte, Marieke Mur, and Peter A Bاندettini. 2008. Representational similarity analysis—connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2:4.
- Liunian Harold Li, Mark Yatskar, Da Yin, Cho-Jui Hsieh, and Kai-Wei Chang. 2019. Visualbert: A simple and performant baseline for vision and language. *arXiv preprint arXiv:1908.03557*.
- Delwin T Lindsey and Angela M Brown. 2014. The color lexicon of american english. *Journal of vision*, 14(2):17–17.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- John Locke. 1847. *An essay concerning human understanding*. Kay & Troutman.
- Li Lucy and Jon Gauthier. 2017. Are distributional representations ready for the real world? evaluating word vectors for grounded perceptual meaning. *arXiv preprint arXiv:1705.11168*.
- Ken McRae, George S Cree, Mark S Seidenberg, and Chris McNorgan. 2005. Semantic feature production norms for a large set of living and nonliving things. *Behavior research methods*, 37(4):547–559.
- William Merrill, Yoav Goldberg, Roy Schwartz, and Noah A Smith. 2021. Provable limitations of acquiring meaning from ungrounded form: What will future language models understand? *arXiv preprint arXiv:2104.10809*.
- Joakim Nivre, Marie-Catherine de Marneffe, Filip Ginter, Jan Hajič, Christopher D. Manning, Sampo Pyysalo, Sebastian Schuster, Francis Tyers, and Daniel Zeman. 2020. **Universal Dependencies v2: An evergrowing multilingual treebank collection**. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4034–4043, Marseille, France. European Language Resources Association.
- Fabio Petroni, Tim Rocktäschel, Sebastian Riedel, Patrick Lewis, Anton Bakhtin, Yuxiang Wu, and Alexander Miller. 2019. **Language models as knowledge bases?** In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2463–2473, Hong Kong, China. Association for Computational Linguistics.
- Tiago Pimentel, Naomi Saphra, Adina Williams, and Ryan Cotterell. 2020. Pareto probing: Trading off accuracy for complexity. *arXiv preprint arXiv:2010.02180*.
- Edoardo Provenzi. 2020. Geometry of color perception. part 1: structures and metrics of a homogeneous color space. *The Journal of Mathematical Neuroscience*, 10(1):1–19.
- Terry Regier, Paul Kay, and Naveen Khetarpal. 2007. Color naming reflects optimal partitions of color space. *Proceedings of the National Academy of Sciences*, 104(4):1436–1441.
- Emily Reif, Ann Yuan, Martin Wattenberg, Fernanda B Viegas, Andy Coenen, Adam Pearce, and Been Kim. 2019a. **Visualizing and measuring the geometry of bert**. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Emily Reif, Ann Yuan, Martin Wattenberg, Fernanda B Viegas, Andy Coenen, Adam Pearce, and Been Kim. 2019b. Visualizing and measuring the geometry of bert. *Advances in Neural Information Processing Systems*, 32:8594–8603.
- Karen J Rossiter. 1996. Structure-odor relationships. *Chemical reviews*, 96(8):3201–3240.
- Dana Rubinstein, Effi Levi, Roy Schwartz, and Ari Rappoport. 2015. How well do distributional models capture different types of semantic knowledge? In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 726–730.
- Armin Saysani, Michael C Corballis, and Paul M Corballis. 2018. Colour envisioned: Concepts of colour in the blind and sighted. *Visual Cognition*, 26(5):382–392.
- Hinrich Schütze. 1992. Dimensions of meaning. In *SC*, pages 787–796.
- Vered Shwartz and Yejin Choi. 2020. **Do neural language models overcome reporting bias?** In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6863–6870, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Milan Straka, Jan Hajic, and Jana Straková. 2016. Udpipeline: trainable pipeline for processing conll-u files performing tokenization, morphological analysis, pos tagging and parsing. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC’16)*, pages 4290–4297.
- Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. 2019. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7464–7473.

- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. [BERT rediscovers the classical NLP pipeline](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4593–4601, Florence, Italy. Association for Computational Linguistics.
- Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Well-read students learn better: On the importance of pre-training compact models. *arXiv preprint arXiv:1908.08962*.
- Ivan Vulić, Edoardo Maria Ponti, Robert Litschko, Goran Glavaš, and Anna Korhonen. 2020. Probing pretrained language models for lexical semantics. *arXiv preprint arXiv:2010.05731*.
- Nathaniel Weir, Adam Poliak, and Benjamin Van Durme. 2020. [Probing neural language models for human tacit assumptions](#).
- Yi Yang, Zhigang Ma, Feiping Nie, Xiaojun Chang, and Alexander G Hauptmann. 2015. Multi-class active learning by uncertainty sampling with diversity maximization. *International Journal of Computer Vision*, 113(2):113–127.
- Noga Zaslavsky, Charles Kemp, Terry Regier, and Naf-tali Tishby. 2018. Efficient compression in color naming and its evolution. *Proceedings of the National Academy of Sciences*, 115(31):7937–7942.

A List of included color terms

Red, green, maroon, brown, black, blue, purple, orange, pink, yellow, peach, white, gray, olive, turquoise, violet, lavender, and aqua.

B RSA between models

Figure 5 shows a the result of representation similarity analysis between the representations derived from all models (and configurations) as well as CIELAB, showing Kendall’s correlation coefficient between flattened RSMs.

C Representation Similarity Matrices

Figures 6 to 9 show the representation similarity matrices employed for the RSA analyses, for the layer with the highest RSA score from each of the controlled-context (CC) models.

D Warm vs. Cool colors

Figures 10 and 11 show Linear Mapping and RSA results broken down by color temperature. The color space is split according to temperature measured according to the Hue dimension in the Hue-Value-Saturation space¹³.

E Corpus statistics

Figures 12 and 13 show log frequency and entropy of distributions over part-of-speech categories, dependency relations, and lemmas of dependency tree heads of color terms in common crawl.

F Linear mapping results by munsell color chip

Figure 14 shows linear mapping results broken down by Munsell chip for all models and configurations.

G Linear mapping control task and probe complexity

Figure 15 shows the full results over a range of probe complexities for the standard experimental condition as well the random control task.

H Dimensionality of color subspace

Figure 16 shows the proportion of explained variance with respect to the number of dimensions which are assigned 95% of the linear regression coefficient weights.

I Effect of model size

Our model size experiments are run using four BERT models of different sizes: BERT-mini (4 layers, hidden size: 256), BERT-small (4 layers, hidden size: 512), BERT-medium (8 layers, hidden size: 512), and BERT-base (12 layers, hidden size: 768). Further model specification and training details for the first three can be found in [Turc et al. \(2019\)](#) and for last in [Devlin et al. \(2019\)](#).

J Linear Mixed Effects Model

To fit Linear Mixed Effects Models, we use the LME4 package. With model type (BERT-CC, RoBERTa-NC, etc.) as a random effect, we follow a step-wise model construction sequence which proceeds along four levels of nesting: (i) in the first level color log-frequency is the only fixed effect, (ii) in the second `pmi-colloc` is added to that, (iii) in the third, each of `pos-ent`, `deprel-ent`, `head-ent` is added separately to the a model with log frequency and `pmi-colloc`, (iv) the term that leads to the best fit from the previous level `deprel-ent` is included, then each of the proportion terms `adj-prop`, `amod-prop`, `cop-prop` is added. The reported regression coefficients are extracted from the minimal model containing each term.

¹³https://psychology.wikia.org/wiki/HSV_color_space

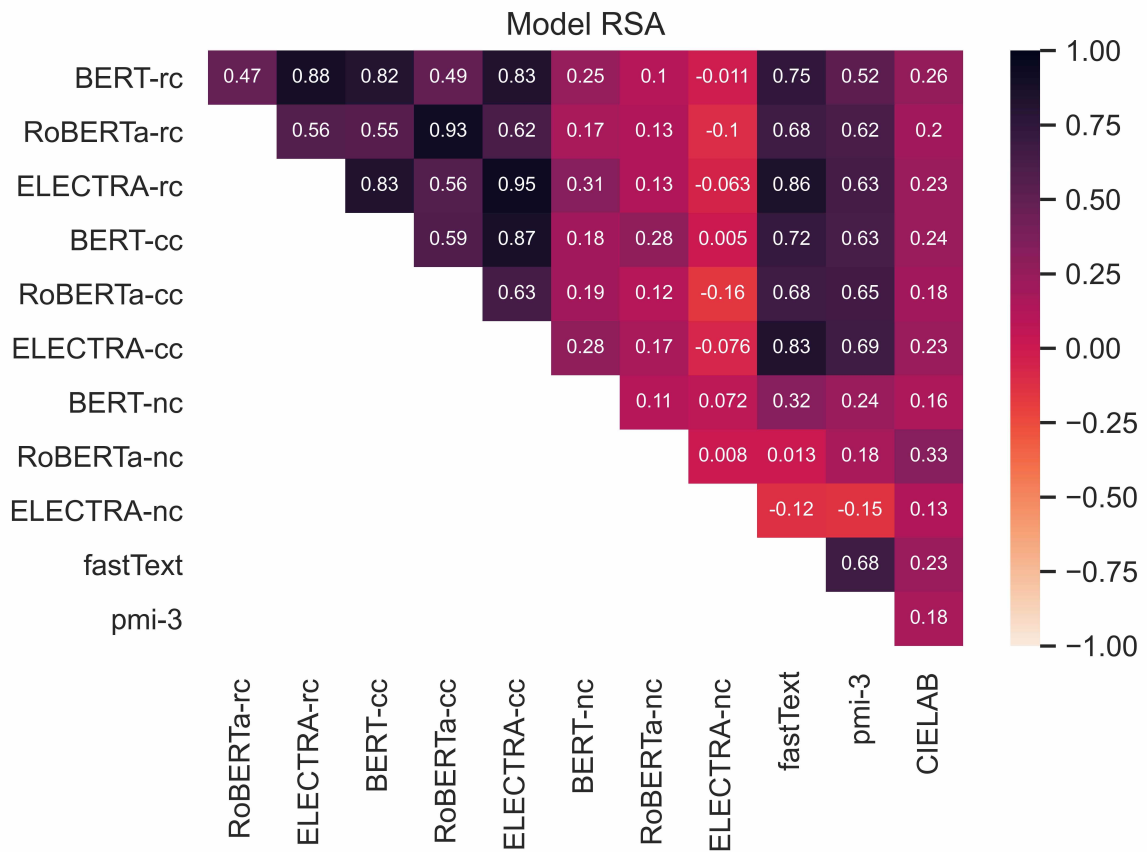


Figure 5: Result of representation similarity analysis between all models (and configurations), showing Kendall's correlation coefficient between flattened RSMs. Results are shown for layers which are maximally correlated with CIELAB, per model. `-rc` indicates **random-context**, `-cc` indicates **controlled-context**, and `-nc` indicates **non-context**.




Figure 6: CIELAB RSM

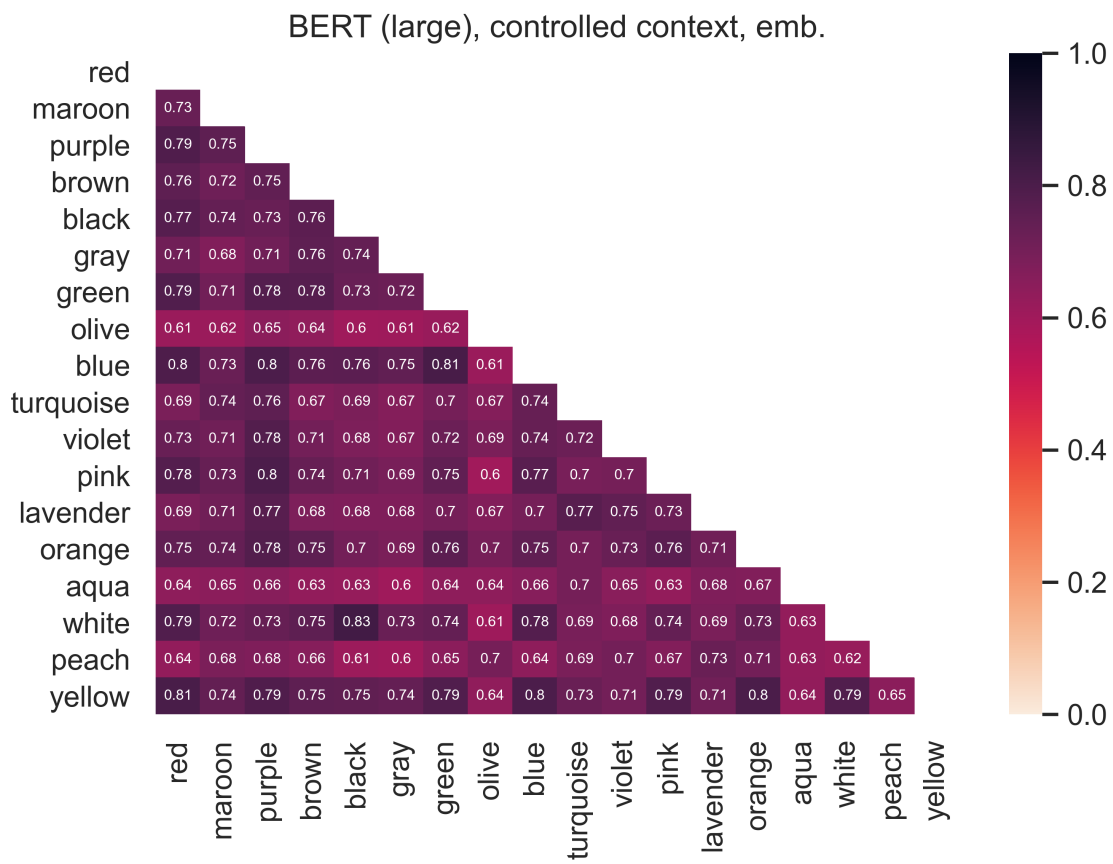


Figure 7: BERT(CC) RSM

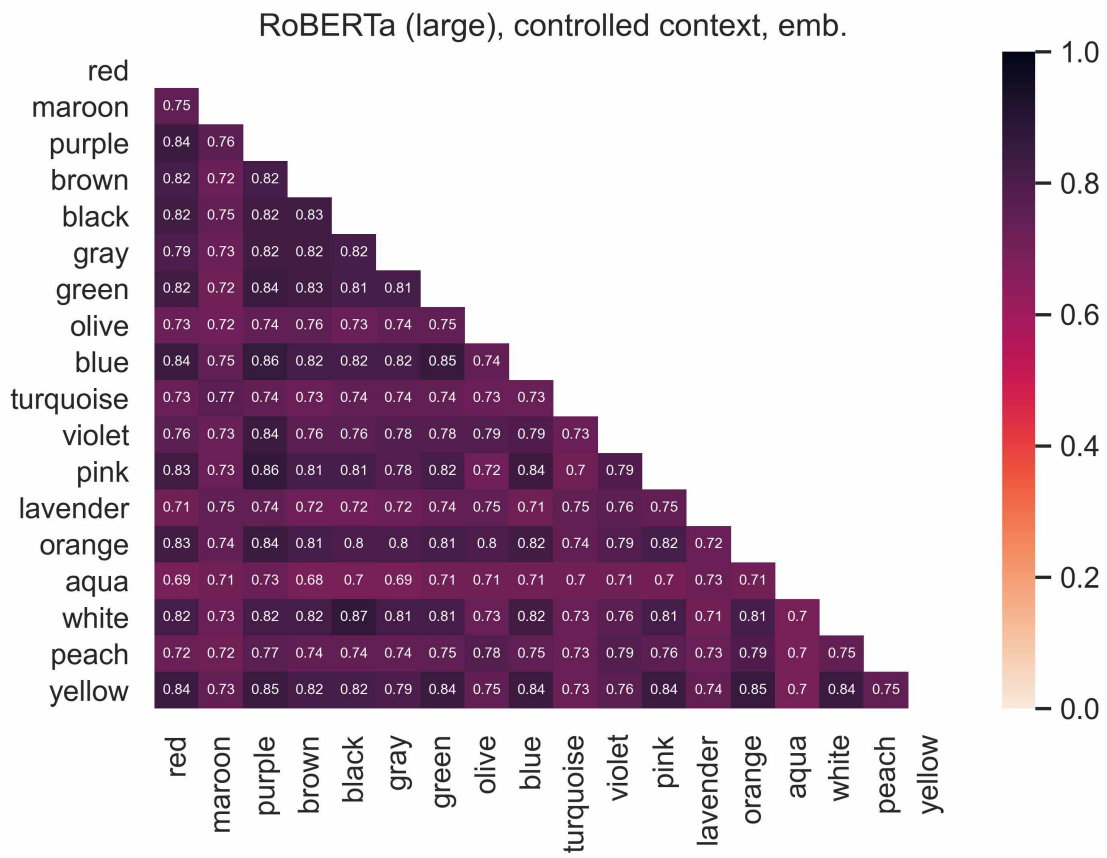


Figure 8: RoBERTa(CC) RSM




Figure 9: ELECTRA(CC) RSM

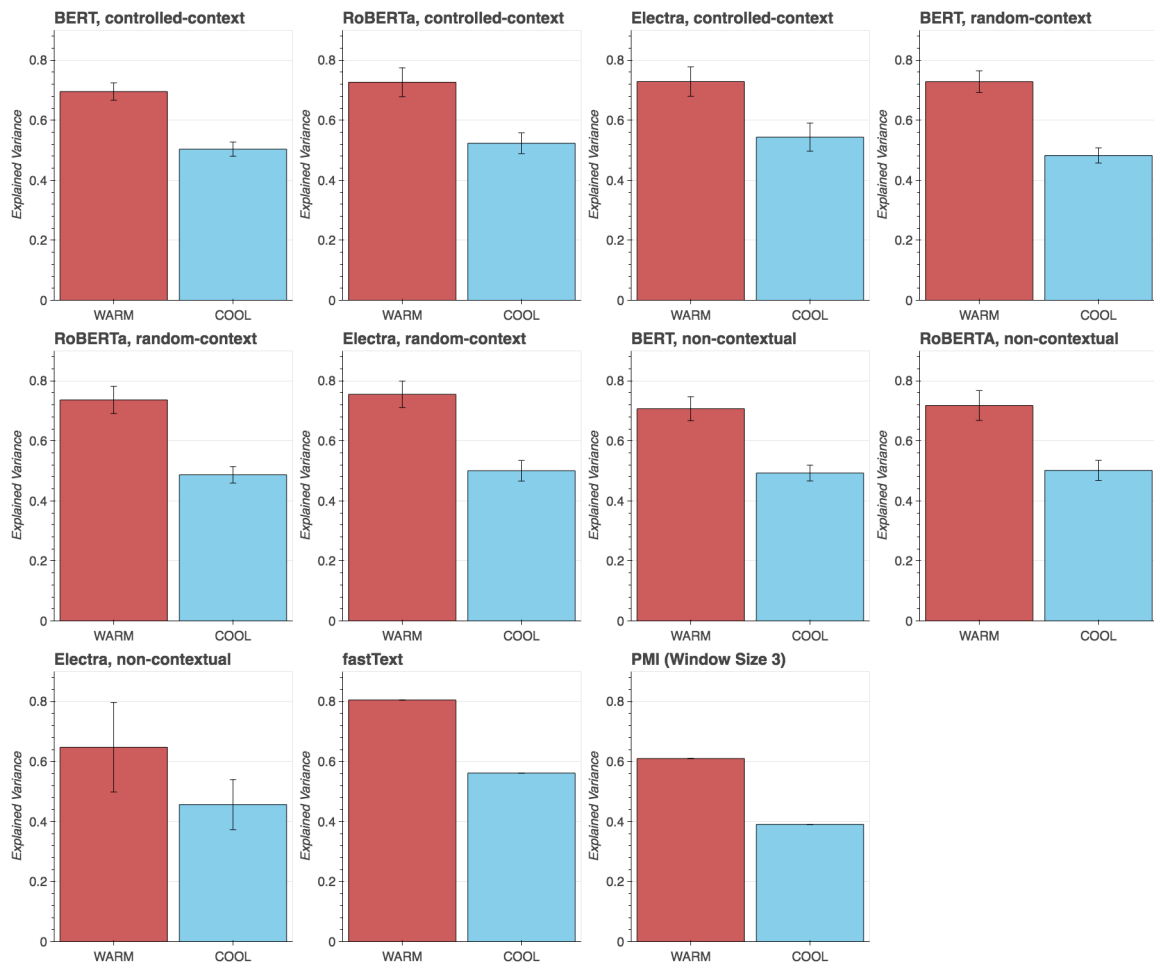


Figure 10: Linear mapping results (proportion of explained variance) broken down by color chip temperature for each of the baselines and the LMs.

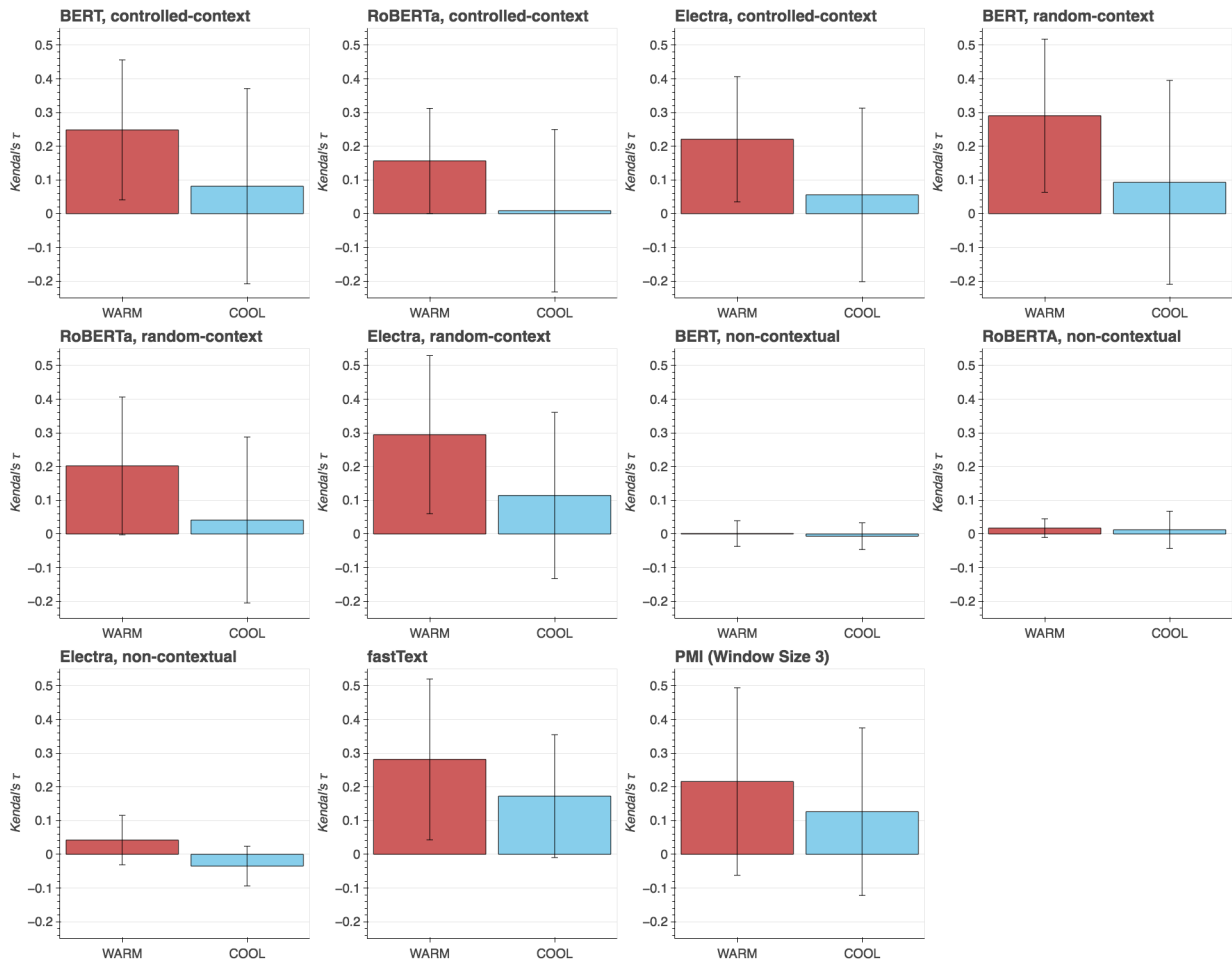


Figure 11: RSA results (Kendall's τ) broken down by color temperature for each of the baselines and the LMs.




Figure 12: Log frequency of color terms in common crawl.




Figure 13: Entropy of distributions over part-of-speech categories, dependency relations, and lemmas of dependency tree heads of color terms in common crawl.

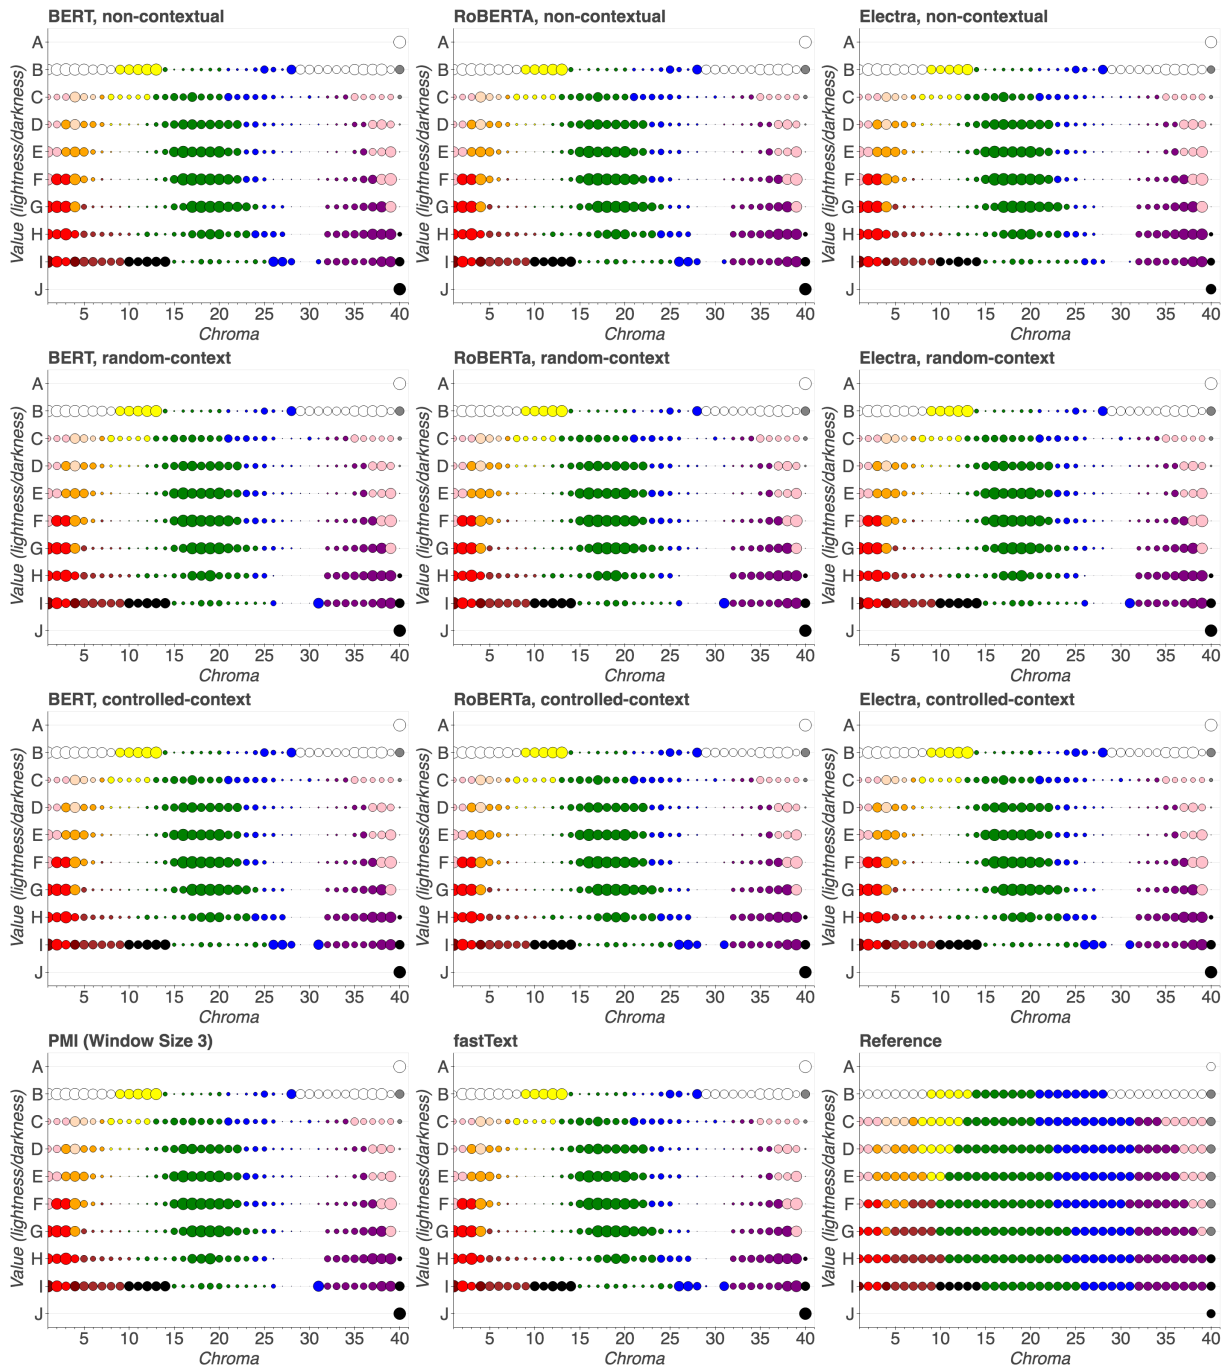


Figure 14: Linear mapping results for each of the baselines and language models, under all extraction configurations, broken down by Munsell color chip. Each circle on the chart represents the ranking of the predicted color chip when ranked according to Pearson distance ($1 - \text{Pearson's } r$) from gold – the larger the circle, the higher (better) the ranking. Circle colors reflect the modal color term assigned to the chips in the lexicon. Reference plot showing modal color of all chips also included.

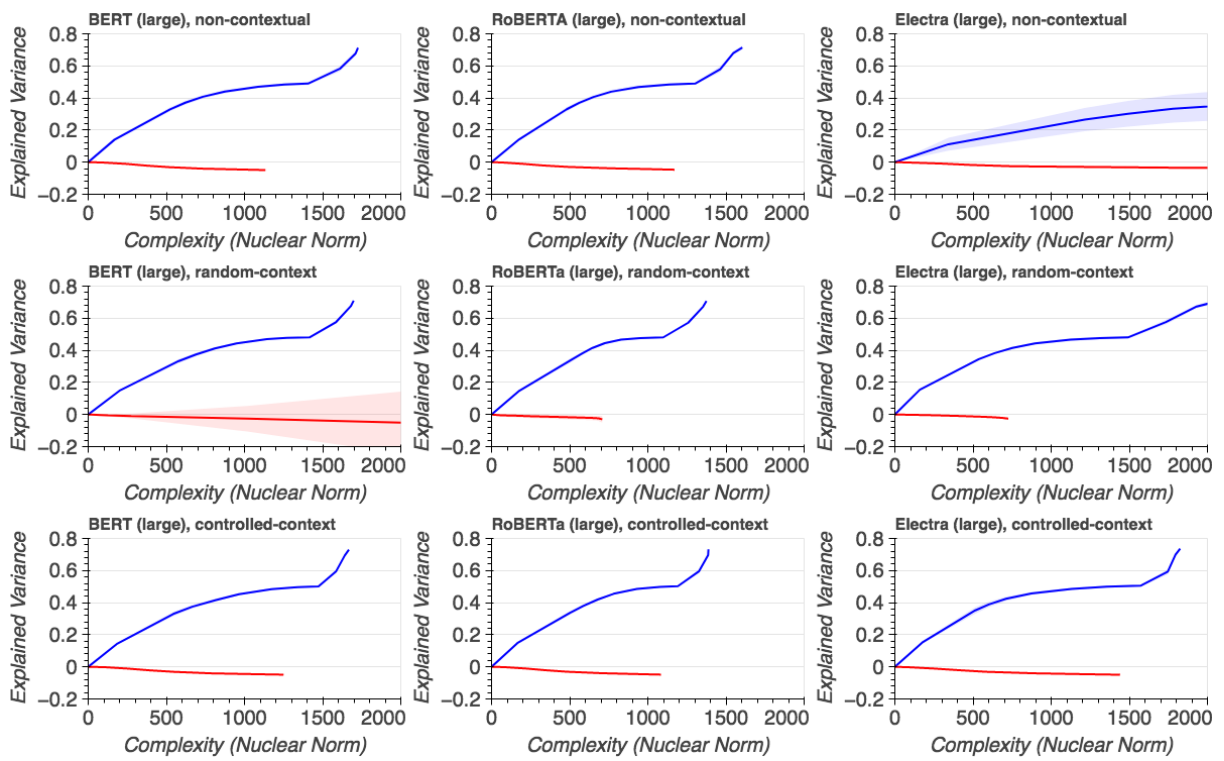


Figure 15: Explained variance for the linear probes trained on the normal experimental condition (blue) and the control task (red) where color terms are randomly permuted. The means are indicated by the lines and standard deviation across layers is indicated by the bands.

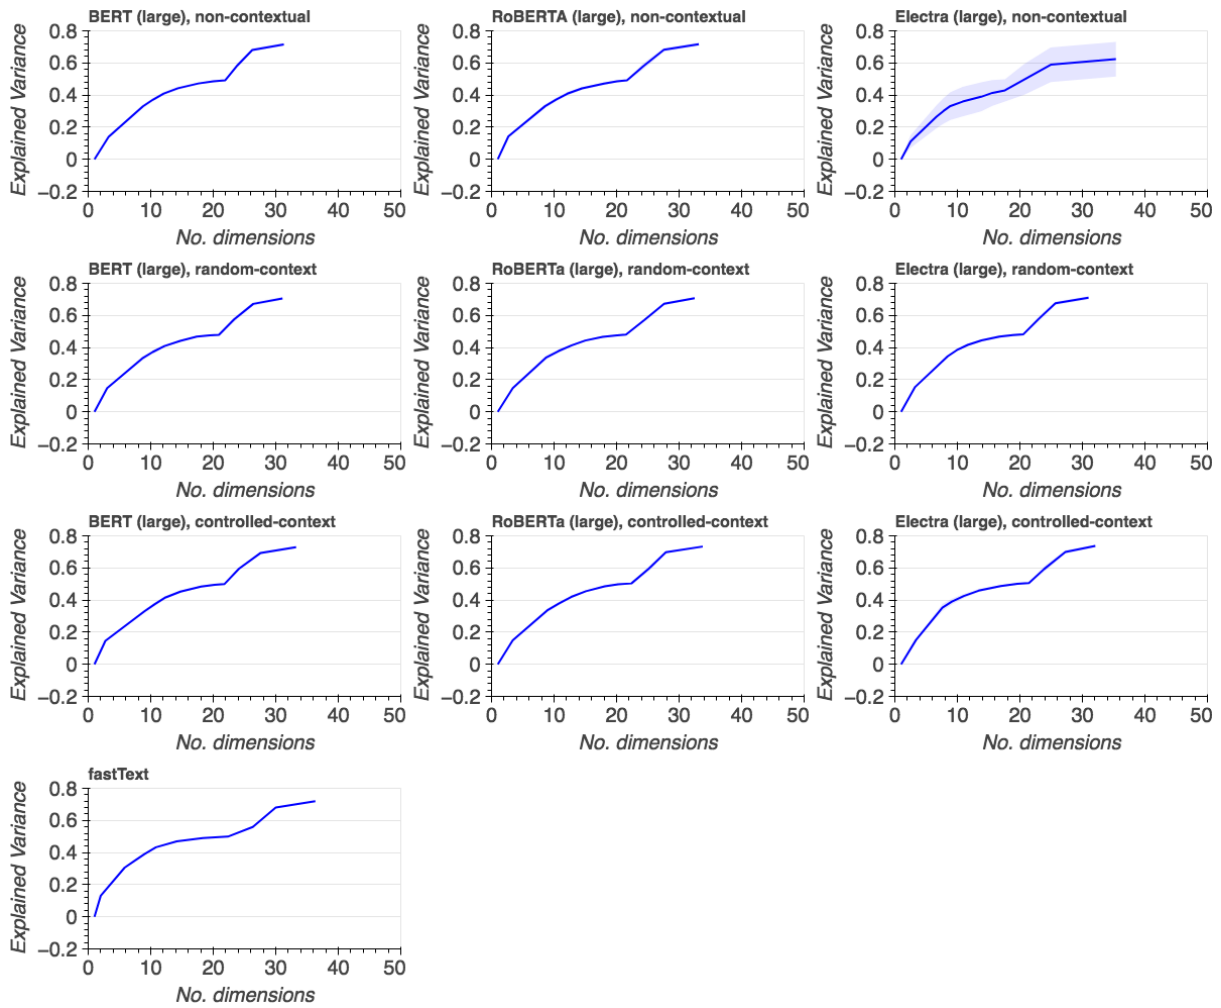


Figure 16: The y-axis shows explained variance for the linear probes. The means are indicated by the lines and standard deviation across layers is indicated by the bands. The x-axis shows the number of regression matrix coefficients assigned 95% of the weight.