

# Interpretable Semantic Gradients in SSD: A PCA Sweep Approach and a Case Study on AI Discourse

**Hubert Plisiecki**  
IDEAS Research Institute  
hplisiecki@gmail.com

**Maria Leniarska**  
VIZJA University  
maria.leniarska@gmail.com

**Jan Piotrowski**  
Warsaw University of Technology  
jan.piotrowski@pw.edu.pl

**Marcin Zajenkowski**  
University of Warsaw  
marcin.zajenkowski@gmail.com

## Abstract

Supervised Semantic Differential (SSD) is a mixed quantitative–interpretive method that models how text meaning varies with continuous individual-difference variables by estimating a semantic gradient in an embedding space and interpreting its poles through clustering and text retrieval. SSD applies PCA before regression, but currently no systematic method exists for choosing the number of retained components, introducing avoidable researcher degrees of freedom in the analysis pipeline. We propose a *PCA sweep* procedure that treats dimensionality selection as a joint criterion over representation capacity, gradient interpretability, and stability across nearby values of  $K$ . We illustrate the method on a corpus of short posts about artificial intelligence written by Prolific participants who also completed Admiration and Rivalry narcissism scales. The sweep yields a stable, interpretable Admiration-related gradient contrasting optimistic, collaborative framings of AI with distrustful and derisive discourse, while no robust alignment emerges for Rivalry. We also show that a counterfactual using a high-PCA dimension solution heuristic produces diffuse, weakly structured clusters instead, reinforcing the value of the sweep-based choice of  $K$ . The case study shows how the PCA sweep constrains researcher degrees of freedom while preserving SSD’s interpretive aims, supporting transparent and psychologically meaningful analyses of connotative meaning.

## 1 Introduction

Supervised Semantic Differential (SSD) is a recent mixed-method that analyzes shifts in meaning of concepts and texts with regard to psychological individual differences of their authors (Plisiecki et al., 2026). The method is inspired by longstanding psychological Semantic Differential methods that measure connotative meaning of concepts via Likert-scales with polar concept opposites (e.g. warm/cold, strong/weak; Osgood et al., 1978) and

by modern distributional semantics, where word embeddings model relational structure in language use (Mikolov et al., 2013; Kozłowski et al., 2019; Garg et al., 2018). Given texts annotated with an outcome variable, such as a trait score or attitude scale, SSD creates individual representations of specific concepts, also called Personal Concept Vectors (PCVs) by aggregating words that collocate with a lexicon of interest relating to a concept under analysis. SSD then fits a linear model in a reduced embedding space with regard to the outcome variable and uses the regression coefficients as a semantic direction or gradient, whose positive and negative poles are explored through clustering of nearby words and further elucidated by retrieving original text snippets which numerical representations are the most aligned to each cluster centroid. The statistical results of the regression are analyzed to estimate whether a linear shift in individual semantic representation explains a significant portion of the variance in the outcome variable.

In the original SSD method, the PCVs are compressed with principal component analysis (PCA) before regression in order to make it possible to apply it to small corpora, and reduce the number of redundant dimensions (Jolliffe, 2002). PCA is the appropriate dimensionality-reduction method for SSD as its linear transformation allows the recovered gradient to be back-projected into the original embedding space (Plisiecki et al., 2026). However, as of now there is no systematic method of choosing the number of components ( $K$ ) to be extracted. While original results were relatively stable across similar component ranges, leaving the choice of dimensionality to the researcher introduces an avoidable researcher degree of freedom, thereby increasing the risk of overfitting, reducing the transparency of the analysis pipeline, and potentially biasing substantive interpretations of the resulting semantic gradients (Simmons et al., 2011).

This work proposes a simple computational heuristic that can be used to pick  $K$  in a principled way. We introduce a PCA sweep procedure that treats  $K$  selection as a joint optimization problem over representation quality, interpretability, and the stability of the results. The sweep evaluates a sequence of  $K$  values, fits SSD at each  $K$  and tracks both cluster coherence (used as an interpretability criterion; Mimno et al., 2011), their alignment with the gradient, as well as a result stability measure based on cosine differences between consecutive gradients. Beyond restricting researcher degrees of freedom, the PCA sweep therefore helps identify a dimensionality at which the recovered semantic gradient remains both stable and interpretable, where interpretability is understood as the extent to which the gradient organizes semantically coherent local neighborhoods that can be meaningfully characterized through nearest neighbors and representative text snippets.

We showcase the applicability of SSD with a PCA sweep by applying it to a dataset of posts written by Prolific study participants who have also answered a questionnaire assessing their Admiration (ADM) and Rivalry (RIV) Narcissistic traits (Back, 2018). The case study serves primarily as an illustration of the extended method, demonstrating how the sweep can constrain dimensionality selection and support more transparent interpretation of recovered semantic gradients under realistic small-sample conditions. The paper’s primary contribution is a methodological procedure for more principled dimensionality selection in SSD, supporting its use as a transparent tool for studying psychologically and socially meaningful variation in language.

## 2 Background: Supervised Semantic Differential

Supervised Semantic Differential (SSD) assumes a collection of documents  $d_i$  paired with continuous outcomes  $y_i$ , such as trait or attitude scores. In its canonical formulation (Plisiecki et al., 2026), SSD focuses on a specific concept of interest by means of a lexicon: words that denote or closely relate to the target concept (e.g., *abortion*, *immigration*). These lexicon terms are used to identify and aggregate the local semantic contexts in which the concept appears, yielding a PCV for each author, reflecting how the concept is represented in their texts.

Each document is mapped to a dense vector  $\mathbf{x}_i \in R^D$  using a fixed embedding model and a composition scheme (e.g., SIF-weighted averages with removal of the top principal component to reduce anisotropy; Arora et al., 2017; Mu et al., 2017). To reduce redundancy and to make the regression applicable in small datasets, SSD applies PCA to the vectors and projects them to a lower-dimensional representation  $\tilde{\mathbf{x}}_i \in R^K$ . A linear model is then estimated:

$$y_i = \alpha + \beta^\top \tilde{\mathbf{x}}_i + \epsilon_i.$$

The regression coefficients are normalized to unit length to obtain a semantic gradient  $\hat{\beta}$ , which is back-projected to the original embedding space. SSD explains the positive and negative poles of this gradient by retrieving nearest neighbors, clustering them separately on each side, and retrieving original text snippets aligned with cluster centroids. These clusters provide qualitative evidence of how the meaning of the focal concept shifts as the outcome increases.

The choice of PCA dimensionality  $K$  affects both regression and interpretation: too few components may collapse distinct semantic regularities, whereas too many may introduce unstable or noisy directions. A principled way to choose  $K$  is therefore essential.

## 3 PCA Sweep for SSD

We introduce a *PCA sweep* procedure that operationalizes the choice of  $K$  as a structured model-selection problem. The key principle is that  $K$  is not optimized for predictive accuracy, but for three properties central to SSD as an interpretive method: (i) the quality of the semantic representation after dimensionality reduction, (ii) the interpretability of the derived semantic gradient, and (iii) the stability of that gradient across nearby values of  $K$ . Although these diagnostics are computed after the regression is fitted, they are not functions of variance explained by the regression or predictive fit, and therefore do not reward models that align more strongly with the outcome variable. Instead, they privilege solutions whose semantic structure is coherent and stable across  $K$ .

Given a user-specified range of candidate dimensionalities (e.g.,  $K \in \{20, 22, \dots, 120\}$ ), the sweep runs SSD repeatedly with identical preprocessing, weighting, and clustering settings, and records diagnostic quantities for each  $K$ .

**Representation statistics.** For every configuration, we compute the cumulative proportion of variance explained by the first  $K$  principal components. This provides a monotonic measure of representation capacity: larger  $K$  values encode more information from the PCVs but also increase the risk of capturing idiosyncratic or noisy variation. Because this quantity grows mechanically with  $K$ , it is used only as a baseline for detrending rather than as a selection criterion.

**Interpretability diagnostics.** After fitting the regression and back-projecting the semantic gradient  $\hat{\beta}_K$  to the embedding space, SSD retrieves nearest neighbors on the positive and negative poles and clusters them separately. For each cluster we compute internal coherence and the cosine alignment between its centroid and  $\hat{\beta}_K$ . These are aggregated into a single, cluster size-weighted *interpretability score* that summarizes how well the learned direction organizes semantically meaningful neighborhoods. To account for the fact that higher-dimensional representations tend to yield higher scores even when the underlying structure does not qualitatively improve, the aggregate is detrended with respect to the (log-transformed) variance explained, and the residuals are standardized to  $z$ -scores.

**Stability diagnostics.** The sweep also tracks how the semantic gradient changes as  $K$  increases. Let  $\hat{\beta}_K$  and  $\hat{\beta}_{K-1}$  be consecutive normalized gradients (back-projected to the original space). We measure *gradient change*

$$\Delta_K = 1 - \cos(\hat{\beta}_K, \hat{\beta}_{K-1}),$$

where lower values indicate greater stability. Intuitively, highly unstable regimes signal that the semantic structure depends on components that are weakly supported by the data.

**Plateau-sensitive smoothing.** Because the interpretability is highly volatile at low  $K$ , instead of rewarding isolated local maxima, both the detrended interpretability curve and the stability curve are smoothed using a local neighborhood average (AUCK), which emphasizes broad, stable plateaus over sharp spikes. This favors configurations where interpretability remains high while the gradient has largely converged.

**Selection rule.** For each  $K$ , the sweep computes a joint score combining interpretability and stability:

$$\text{joint\_score}_K = \frac{1}{2}(\text{interp\_auck}_K + \text{stab\_auck}_K).$$

The selected dimensionality is the *smallest*  $K$  attaining the maximal joint score, privileging parsimonious, stable, and interpretable solutions.

## 4 Case Study: Admiration and Rivalry in AI Discourse

We applied SSD with the PCA sweep to a corpus of  $N=349$  short posts about artificial intelligence (AI) written by Prolific participants (mean length  $\approx 30$  words; full writing prompt and additional study details provided in Appendix A), who also completed Admiration (ADM) and Rivalry (RIV) narcissism scales reflecting the two self-regulatory strategies in the Narcissistic Admiration and Rivalry Concept (Back et al., 2013), where Admiration captures assertive self-enhancement and status-seeking tendencies, and Rivalry reflects defensive, antagonistic self-protection in response to perceived threat (Back, 2018). Tokenization and linguistic preprocessing were performed using spaCy 3.8.7 with the en\_core\_web\_sm English model (Montani et al., 2023). Posts were embedded using the 300 dimensional Dolma GloVe model (Carlson et al., 2025) with SIF weighting ( $a=10^{-3}$ ), removal of the top principal component. Crucially, a lexicon was not used to dial down on AI-specific terms, because the posts did not rely on a consistent set of shared AI keywords; however, since all texts were produced under the same AI-focused prompt, each post can be treated as a whole as a meaningful reflection of participants’ reactions to AI. The sweep evaluated  $K \in \{1, 3, \dots, 119\}$  with cluster number chosen by silhouette from ranges  $k \in [2, 5]$ , top-100 neighbors per pole, AUCK radius = 3, and median-smoothed gradient change (win= 7); remaining settings followed the original SSD configuration<sup>1</sup>.

Table 1 summarizes the regression results at the sweep-selected  $K$  values. For ADM, the model explained a small–moderate but reliable proportion of variance ( $R_{\text{adj}}^2 = .19$ ,  $F = 6.32$ ,  $p < 10^{-10}$ ,  $r \approx .47$ ) with a pronounced semantic gradient ( $\|\hat{\beta}\| = 5.58$ ;  $\Delta_{0.1} = 0.65$ ). By contrast, the RIV model did not reach significance ( $R_{\text{adj}}^2 = .03$ ,  $p = .095$ ), and we therefore refrain from qualitative interpretation of the RIV gradient.

This non-significant RIV result should be interpreted cautiously. Prior SSD power analyses suggest that, under idealized conditions with high-quality on-topic texts and fixed PCA dimensional-

<sup>1</sup>Code and supplementary materials are available at [github.com/hplisiecki/acl2026-ssd-pca-sweep](https://github.com/hplisiecki/acl2026-ssd-pca-sweep)

ity, a corpus of this size would generally be sufficient to detect small effects. Under more realistic conditions involving short texts and mixed text quality, however, required sample sizes can be substantially larger, placing the present dataset in a range that is more clearly suited to detecting medium than small effects (Plisiecki et al., 2026). At the same time, datasets of this scale are typical in psychological research combining open-ended responses with trait measures, making the present corpus appropriate for a PCA sweep case study under realistic small-sample conditions. The null RIV result may therefore reflect either a genuine absence of stable semantic alignment for Rivalry in this corpus or a Type II error due to limited power under realistic conditions. For this reason, we do not interpret the RIV gradient qualitatively, but we also do not treat this null result as evidence against the sweep procedure itself.

| Trait | $K$ | $R_{\text{adj}}^2$ | $F$  | $p$          | $r$ | $\ \hat{\beta}\ $ |
|-------|-----|--------------------|------|--------------|-----|-------------------|
| ADM   | 15  | .19                | 6.32 | $< 10^{-10}$ | .47 | 5.58              |
| RIV   | 23  | .03                | 1.43 | .095         | .30 | 5.38              |

Table 1: SSD regression results for AI posts at the sweep-selected  $K$ .




Figure 1: Sweep diagnostics for ADM: detrended interpretability and smoothed unit-change curves, with the selected  $K=15$  plateau.

Figure 1 presents the sweep diagnostics for ADM. The detrended interpretability curve rises sharply at low  $K$  and reaches a peak around  $K=15$  (explaining 50% of variance in the original PCVs). The unit-change curve exhibits a local plateau at the same position; although it declines for larger  $K$ , this drop is mirrored by a decrease in interpretability. Taken together, these trends indicate that  $K=15$  lies in the most parsimonious region where the gradient remains both interpretable and stable.

To interpret the ADM gradient, we clustered nearest neighbours on the positive and negative poles. Table 2 summarizes the resulting themes with snippets of relevant representative excerpts.

| Pole | Size | Summary (Top Words / Excerpt)   |
|------|------|---|
| +    | 14   | <i>Cultivation &amp; enrichment: cultivate, re-discover, rejuvenated</i> — “AI’s potential is boundless...” |
| +    | 86   | <i>Innovation &amp; collaboration: innovation, partnership, empower</i> — “AI is transforming our world...” |
| –    | 56   | <i>Deception &amp; threat: misleading, dishonest, unfair</i> — “woke, evasive and bizarre responses...”     |
| –    | 44   | <i>Ridicule &amp; contempt: ridiculous, absurd, laughable</i> — “the stupid programmers exposed their hand” |

Table 2: Positive and negative ADM clusters in AI discourse.

The positive pole reflects a *future-oriented, collaborative, and prosocial framing of AI*, linking it to innovation, integration, and constructive technological progress. The negative pole instead emphasizes *distrust, antagonism, and derision*, portraying AI and its creators as deceptive, biased, or ideologically motivated. Overall, the case study illustrates how the PCA sweep supports interpretable SSD analysis by constraining dimensionality in a transparent, stability-aware manner.

For comparison, Appendix B reports a counterfactual SSD run at an arbitrarily high dimensionality ( $K=120$ ), where the resulting clusters become diffuse and harder to interpret. They include topics unrelated to AI such as gardening, and wildlife, and foreign vocabulary, illustrating the usefulness of the PCA Sweep. In order to visualize the contrast between the sweep-selected and high-dimensional solutions, we projected the retrieved lexical neighbors around both poles of the back-projected gradient into two dimensions using PCA fitted separately within each solution. Figure 2 shows that the sweep-selected solution yields a more compact and clearly organized arrangement of neighbors, whereas the high-dimensional counterfactual is more fragmented and diffuse, especially on the negative pole. Although this visualization is only qualitative, it is consistent with the broader diagnostic pattern. To probe robustness beyond a single embedding model, Appendix C also reports a reanalysis of the same dataset, yielding a substantively similar pattern of results.




Figure 2: Local neighborhood structure for ADM in the sweep-selected solution ( $K=15$ ; left) and the high-dimensional counterfactual ( $K=120$ ; right). Data point colors correspond to separate clusters.

## 5 Discussion and Conclusion

This work introduced a PCA sweep procedure for SSD that selects dimensionality by jointly prioritizing gradient stability and semantic interpretability rather than predictive fit. In the AI case study, the sweep yielded a compact representation with coherent cluster structure and stable gradients, whereas the counterfactual high-dimensional run in Appendix B produced diffuse and weakly structured clusters. This contrast demonstrates that the sweep not only constrains researcher degrees of freedom, but also helps preserve the qualitative aims of SSD by discouraging over-parameterized representations that cease to reflect psychologically meaningful structure.

Beyond the methodological contribution, the case study revealed a psychologically interpretable association between Admiration and the semantics of AI discourse. The positive pole of the ADM gradient reflected future-oriented, collaborative, and prosocial framings of AI, while the negative pole emphasized distrust, antagonism, and ridicule. This contrast is consistent with accounts of Admiration as an agentic, self-enhancing orientation toward status and positively valenced self-views (Back, 2018; Zeigler-Hill et al., 2019), suggesting that individuals higher in Admiration may align themselves with culturally valued narratives of technological progress and collective advancement. A complementary interpretation draws on prior work describing large AI systems as compliant and *sycophantic* in evaluative interaction (Perez et al., 2022); semantics that frame AI as supportive, empowering,

or socially harmonious may resonate more strongly with affirmation-oriented self-regulatory motives linked to Admiration.

More broadly, the results illustrate how SSD, when paired with a principled choice of PCA dimensionality, can surface stable, interpretable semantic gradients that connect language use to underlying psychological dispositions. However, the PCA sweep addresses only one source of flexibility in the SSD pipeline. A key open challenge is developing similarly principled criteria for upstream modeling choices - most notably, the selection of the base embedding model. Different embedding spaces encode distinct cultural, temporal, and stylistic regularities, and these choices can meaningfully shape the structure of recovered gradients. Extending the logic of stability- and interpretability-based diagnostics to the level of model selection represents an important next step toward a fully transparent and methodologically grounded SSD workflow.

Overall, the sweep procedure strengthens SSD as a mixed quantitative–interpretive method by tying dimensionality selection directly to gradient stability and semantic coherence. The approach offers a lightweight, model-agnostic criterion that can be integrated into embedding-based analysis in order to both enhance interpretability and reduce researcher degrees of freedom.

## Limitations

Our case study is based on a relatively small, survey-elicited dataset of short AI-related posts from Prolific participants, which limits the gener-

alizability of the observed semantic gradients beyond this population and task context. The analysis also operates at the level of whole-text representations rather than concept-specific lexicon-based PCVs; while justified by the uniform prompting, this choice may blur finer-grained semantic distinctions. Moreover, the PCA sweep addresses only one source of flexibility in SSD, namely the dimensionality selection, while other design choices remain open, such as the choice of the embedding model, and word window size when using a lexicon. As the appendix counterfactual illustrates, different representational assumptions can meaningfully affect interpretability, and a broader framework for principled model and parameter selection in SSD remains an open methodological challenge. Finally, the psychological interpretations offered here are correlational and depend on the alignment between semantic structure and trait scores; future work should test the robustness of these gradients across datasets, languages, and outcomes, as well as test the case study conclusions with targeted laboratory studies.

### Ethical considerations

The present work analyzes text data paired with individual-difference measures, raising issues related to privacy, interpretative uncertainty, and the potential misuse of trait–language associations. Importantly, SSD is not designed as a predictive or profiling technology: the method estimates weak, low-variance semantic gradients that support qualitative interpretation rather than accurate inference about individuals. As such, the approach is poorly suited for identifying, classifying, or predicting participants’ psychological dispositions, and we explicitly discourage such uses. Our analyses are conducted at the group and pattern level, and all interpretations are contingent, correlational, and theoretically guided rather than diagnostic. More broadly, we view SSD as a tool for hypothesis generation and meaning exploration in psychological and social research, not as a mechanism for automated assessment or decision making. The data was confirmed to not include any personally identifying or offensive content. The text of this manuscript was partially polished with the assistance of a Large Language Model; all automatically revised passages were reviewed, validated, and, where necessary, corrected by the authors to ensure accuracy, faithfulness to the data, and conceptual

integrity. Future applications of SSD should continue to prioritize transparency, participant respect, and careful communication of uncertainty.

### Acknowledgments

This work was supported by the Foundation for Polish Science (FNP) grant No. FENG.02.01-IP.05-0058/24, and by the National Science Center (NCN) grant No. 2023/51/B/HS6/0068.

### References

- Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. [A Simple but Tough-to-Beat Baseline for Sentence Embeddings](#).
- Mitja D. Back. 2018. [The Narcissistic Admiration and Rivalry Concept](#). In Anthony D. Hermann, Amy B. Brunell, and Joshua D. Foster, editors, *Handbook of Trait Narcissism: Key Advances, Research Methods, and Controversies*, pages 57–67. Springer International Publishing, Cham.
- Mitja D. Back, Albrecht C. P. Küfner, Michael Dufner, Tanja M. Gerlach, John F. Rauthmann, and Jaap J. A. Denissen. 2013. [Narcissistic admiration and rivalry: Disentangling the bright and dark sides of narcissism](#). *Journal of Personality and Social Psychology*, 105(6):1013–1037.
- Riley Carlson, John Bauer, and Christopher D. Manning. 2025. [A New Pair of GloVes](#). *arXiv preprint*. Version Number: 1.
- Nikhil Garg, Londa Schiebinger, Dan Jurafsky, and James Zou. 2018. [Word embeddings quantify 100 years of gender and ethnic stereotypes](#). *Proceedings of the National Academy of Sciences*, 115(16).
- I. T. Jolliffe. 2002. *Principal Component Analysis*. Springer Series in Statistics. Springer-Verlag, New York.
- Austin C. Kozlowski, Matt Taddy, and James A. Evans. 2019. [The Geometry of Culture: Analyzing the Meanings of Class through Word Embeddings](#). *American Sociological Review*, 84(5):905–949.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. 2013. [Efficient Estimation of Word Representations in Vector Space](#). *arXiv preprint*. Version Number: 3.
- David Mimno, Hanna Wallach, Edmund Talley, Miriam Leenders, and Andrew McCallum. 2011. [Optimizing Semantic Coherence in Topic Models](#). In *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing*, pages 262–272, Edinburgh, Scotland, UK. Association for Computational Linguistics.

Ines Montani, Matthew Honnibal, Adriane Boyd, Sofie Van Landeghem, and Henning Peters. 2023. [explosion/spaCy: v3.7.2: Fixes for APIs and requirements](#).

Jiaqi Mu, Suma Bhat, and Pramod Viswanath. 2017. [All-but-the-Top: Simple and Effective Postprocessing for Word Representations](#). *arXiv preprint*. Version Number: 2.

Charles Egerton Osgood, George J. Suci, and Percy H. Tannenbaum. 1978. *The measurement of meaning*. University of Illinois Press, Urbana-Champaign.

Ethan Perez, Sam Ringer, Kamilè Lukošiušė, Karina Nguyen, Edwin Chen, Scott Heiner, Craig Pettit, Catherine Olsson, Sandipan Kundu, Saurav Kadavath, Andy Jones, Anna Chen, Ben Mann, Brian Israel, Bryan Seethor, Cameron McKinnon, Christopher Olah, Da Yan, Daniela Amodei, and 44 others. 2022. [Discovering Language Model Behaviors with Model-Written Evaluations](#). *arXiv preprint*. Version Number: 1.

Hubert Plisiecki, Paweł Lenartowicz, Artur Pokropek, Kinga Małyska, and Maria Flakus. 2026. [Measuring Individual Differences in Meaning: The Supervised Semantic Differential](#).

Joseph P. Simmons, Leif D. Nelson, and Uri Simonsohn. 2011. [False-Positive Psychology: Undisclosed Flexibility in Data Collection and Analysis Allows Presenting Anything as Significant](#). *Psychological Science*, 22(11):1359–1366.

Virgil Zeigler-Hill, Jennifer K. Vrabel, Gillian A. McCabe, Cheryl A. Cosby, Caitlin K. Traeder, Kelsey A. Hobbs, and Ashton C. Southard. 2019. [Narcissism and the pursuit of status](#). *Journal of Personality*, 87(2):310–327.

## A Prolific Study Details

For transparency and reproducibility, we report the exact instruction shown to participants in the AI writing task. Participants were asked to write a short text in response to the following prompt:

**Prompt:** *Using the box below, please write a post about artificial intelligence. It can be your opinion, a comment for an application of AI or any other thought about AI, you would like to share.*

All posts in the AI corpus were produced in the English language in response to this prompt, and no additional guidance or examples were provided. Standard informed consent was obtained at the beginning of the study, and participants were informed about its aims and the nature of their participation. The study was prepared and conducted in accordance with the guidelines of the relevant

psychological research ethics committee. Participants were recruited through the online crowdsourcing platform Prolific, which provides access to a diverse international pool of adult participants and allows researchers to apply demographic and eligibility pre-screening criteria. In this study, participants were recruited from the United States; the mean age was 39 years, and the gender distribution was approximately 51% female (179), 48% male (167), and <1% identifying as another gender (3). All participants were members of the general Prolific participant pool and were compensated financially for their participation. Compensation was set on a task-time basis and corresponded to an average effective hourly rate of £9.10, calculated from the observed distribution of completion times in the final dataset. We report the hourly-equivalent rate to provide a transparent indicator of remuneration relative to time on task. The compensation level was consistent with Prolific’s fair-pay norms and with ethical standards for online research at the time of data collection, and was higher than the prevailing federal minimum-wage benchmark in the United States during the same period.

## B SSD at High Dimensionality

For comparison with the PCA sweep results in the main body of the work, we reran SSD using an arbitrarily high dimensionality ( $K=120$ ), corresponding to a near-full-variance reconstruction of the embedding space (96% of the variance explained) rather than a parsimonious, stability-informed representation. All other preprocessing, regression, and clustering settings were held constant.

| No. | Size | Summary (Top Words / Excerpt)   |
|-----|------|---|
| 1   | 51   | <i>Mixed institutional / geographic terms: inclusive, asean, maldives, fiji, rolex, brics, lagos, ap ec</i> — “AI is going to take over the world.”     |
| 2   | 49   | <i>Multilingual / function-word cluster: como, del, para, con, mejor, su, sistema, mundo, uso</i> — “I really love AI; they are very useful to humans.” |

Table 3: High-ADM clusters at  $K=120$ .

Tables 3 and 4 report the clusters obtained at  $K=120$  for the ADM gradient. In contrast to the sweep-selected solution, the clusters at this dimensionality exhibit lower thematic coherence and weaker alignment between lexical content and the semantic gradient, despite the SSD regression itself explaining even more of the variance than the  $K=15$  one (adjusted  $R^2 = 0.234$ ,  $F = 1.89$ ,

$p = 2.17 \times 10^{-5}$ ). Many clusters mix unrelated lexical fields (e.g., geopolitical terms, brand names, animal categories, gardening vocabulary) or aggregate language unrelated to AI discourse, indicating that the gradient begins to track idiosyncratic lexical variation rather than psychologically meaningful semantic organization. This pattern is consistent with the interpretability–stability diagnostics: when the representation is expanded to include high-variance but weakly structured components, the resulting neighborhoods become diffuse and less informative for qualitative interpretation.

| No. | Size | Summary (Top Words / Excerpt)   |
|-----|------|---|
| 1   | 27   | <i>Animal / wildlife lexicon: birds, rabbits, squirrels, pigeons, sparrows, menagerie</i> — “AI is going to kill jobs.”                                   |
| 2   | 31   | <i>Gardening / foliage terms: foliage, bark, potted, mulch, thyme, manure, thistle</i> — “AI is a great tool for problems that have already been solved.” |
| 3   | 11   | <i>Physical-action / nautical fragment: raking, hammering, trampling, keel, broadside</i> — “AI just blows me away!”                                      |
| 4   | 17   | <i>Speculation / rumor vocabulary: alluding, unsubstantiated, tidbit, fable, sidenote</i> — “The stupid programmers exposed their hand...”                |
| 5   | 14   | <i>Pet / illustration terms: terrier, labrador, postcards, ephemera</i> — “Used an AI service to make some cute cat pictures.”                            |

Table 4: Low-ADM clusters at  $K=120$ .

Overall, the  $K=120$  solution demonstrates that maximizing variance alone leads to semantically diffuse clusters that are difficult to interpret in relation to ADM. In contrast, the sweep-selected configuration produces more coherent and psychologically meaningful gradients, underscoring the value of dimensionality choice as an interpretability-relevant decision rather than a purely statistical one.

## C Reanalysis with the Common Crawl 840B Model

To probe the robustness of the main case-study conclusions to the choice of embedding model, we re-ran the analysis on the same AI corpus using the Common Crawl 840B GloVe model (Carlson et al., 2025). All preprocessing and PCA sweep settings were kept identical to the main analysis. As in the main case study, the sweep evaluated  $K \in \{1, 3, \dots, 119\}$  and selected dimensionality by jointly considering representation quality, interpretability, and gradient stability.

Table 5 summarizes the regression results at the

sweep-selected  $K$  values. For ADM, the reanalysis produced a pattern closely comparable to the main Dolma-based result, with a reliable effect of similar magnitude ( $R_{\text{adj}}^2 = .16$ ,  $F = 7.19$ ,  $p < 10^{-10}$ ,  $r \approx .44$ ), although at a slightly lower selected dimensionality ( $K=11$ ). By contrast, the RIV model again did not reach significance ( $R_{\text{adj}}^2 = .01$ ,  $p = .319$ ), despite a larger selected dimensionality ( $K=35$ ). This replication of the null RIV result across embedding models is consistent with the interpretation that Rivalry does not yield a stable and recoverable semantic gradient in this corpus, although limited power for weaker effects cannot be ruled out.

| Trait | $K$ | $R_{\text{adj}}^2$ | $F$  | $p$          | $r$ | $\ \hat{\beta}\ $ |
|-------|-----|--------------------|------|--------------|-----|-------------------|
| ADM   | 11  | .16                | 7.19 | $< 10^{-10}$ | .44 | 3.99              |
| RIV   | 35  | .01                | 1.11 | .319         | .33 | 5.46              |

Table 5: SSD regression results for AI posts with the Common Crawl 840B model at the sweep-selected  $K$ .

| Pole | Size | Summary (Top Words / Excerpt)   |
|------|------|---|
| +    | 32   | <i>Cultivation &amp; empowerment: strive, nurture, opportunities, inspire, empower</i> — “AI is transforming our world in fascinating ways...”                      |
| +    | 17   | <i>Innovation &amp; advancement: innovation, advancements, breakthroughs, pioneering, transformative</i> — “AI is transforming our world in incredible ways...”     |
| +    | 27   | <i>Collaboration &amp; integration: partnerships, collaborative, integrate, technologies, sustainability</i> — “the possibilities are endless...”                   |
| +    | 12   | <i>Excellence &amp; boundlessness: unparalleled, world-class, cutting-edge, limitless, boundless</i> — “the potential is limitless...”                              |
| +    | 12   | <i>Learning &amp; participation: learning, interactive, literacy, learners, mentoring</i> — “AI will revolutionize interactive media.”                              |
| –    | 46   | <i>Error &amp; misleadingness: mistaken, misleading, incorrectly, inaccurate, untrue</i> — “I have to be very careful to fact check the information it gives me...” |
| –    | 54   | <i>Awfulness &amp; irritation: awful, weird, stupid, horrible, annoying, scary</i> — “AI creates really ugly weird art...”  |

Table 6: Positive and negative ADM clusters in AI discourse with the Common Crawl 840B model at the sweep-selected  $K=11$ .

To interpret the ADM gradient, we clustered nearest neighbours on the positive and negative poles. Table 6 summarizes the resulting themes with representative excerpts. The positive pole again reflects a broadly *future-oriented, aspirational, and collaborative framing of AI*, linking it to innovation, empowerment, integration, and

technological progress. The negative pole instead emphasizes *error, unpleasantness, and distrust*, especially around incorrectness, misleadingness, and irritation. Thus, despite differences in selected  $K$  and local neighborhood composition, the broader semantic contrast for ADM remained stable across embedding models.

For comparison, we also inspected a counterfactual SSD run at an arbitrarily high dimensionality ( $K=120$ ). As in the main analysis, the resulting neighborhoods became less semantically consolidated. Although one positive cluster retained an aspirational vocabulary, other clusters became weaker, noisier, or less clearly related to AI discourse. Table 7 summarizes the cluster content retrieved for the ADM counterfactual.

| Pole | Size | Summary (Top Words / Excerpt)  |
|------|------|--|
| +    | 78   | <i>Excellence &amp; prestige: excellence, commitment, unparalleled, world-class, unsurpassed</i> — “AI is transforming industries at an unprecedented pace...”                                   |
| +    | 22   | <i>Multilingual / noisy tokens: na, po, ko, ni, ka, ke, mas, ang, ba, tak</i> — “chat gpt hard asl helps me write dum shi”   |
| –    | 37   | <i>Liquids / gardening spillover: pouring, spraying, watering, puddle, mustard, glaze</i> — “I’m so tired of everything being replaced by this AI hot garbage...”                                |
| –    | 39   | <i>Art objects &amp; decoration: sculptures, paintings, drawings, murals, figurines, posters</i> — “I’m always looking for new ways to use AI in my daily life. Most recent use: garden design.” |
| –    | 24   | <i>Animals &amp; wildlife: deer, rabbits, squirrel, birds, coyote, moose, wolves</i> — “Used an AI service to make some cute and funny cat pictures!”  |

Table 7: Positive and negative ADM clusters in AI discourse with the Common Crawl 840B model at the high-dimensional counterfactual  $K=120$ .

Compared with the sweep-selected solution, the high-dimensional counterfactual is qualitatively less coherent as a basis for interpretation. In particular, several negative-pole neighborhoods drift toward semantically weak or off-topic content, including liquids, decorative objects, and animals, while one positive-pole cluster becomes dominated by multilingual or otherwise noisy tokens. This mirrors the pattern observed in the main analysis and further supports the claim that unconstrained dimensionality can produce diffuse or substantively irrelevant local neighborhoods even when some broad aspirational vocabulary remains recoverable.

Overall, the Common Crawl 840B reanalysis supports the robustness of the main methodolog-

ical conclusion. Across two embedding models, the PCA sweep selected comparatively low-dimensional ADM solutions with coherent substantive structure, while RIV again failed to yield a reliable gradient and the high-dimensional counterfactual produced weaker and less clearly interpretable neighborhoods.