The Gendered Linguistic Geometry of Everyday Spaces in American English Discourse, 1900–2009

Domain-Conditional Feminization and the Asymmetric Emergence of Private Vocabulary

Anonymous submission for double-blind review

Abstract

Sociological accounts of gendered space treat the public-private divide as a structural feature of modern life, but the linguistic expression of this divide is itself historically uneven. We construct decade-aligned word embeddings on Google Books American English (1900–2009, 11 decades, 127 billion words) and apply a five-stage pre-registered vocabulary protocol that anchors candidate space tokens in WordNet hyponym walks from physical-space root synsets, filters them through seven documented exclusion lists, and partitions them by data-driven frequency thresholds. The protocol yields 87 main and 78 supplementary tokens organized into 11 theoretical domains drawn from Spain (1992) and Massey (1994). We then estimate within-space fixed-effects regressions of each cultural axis on the gender axis across six axes (prestige, potency, income, education, morality, activity). At the pooled level, only the potency axis shows a marginally significant effect. Per-domain analysis shows that the pooled prestige null is the average of opposing effects. Workplace spaces feminize alongside prestige decline, Healthcare spaces feminize alongside prestige rise, and on the potency axis four of five significant domains show feminization predicting potency decline, with Workplace the single positive exception. Tokens that first enter the formal lexicon at or after 1950, including kitchen, bedroom, classroom, and restaurant, feminize and lose prestige post-entry at high rates, with Healthcare again the exception. The gendered cultural geometry of American spaces is domain-conditional rather than universal.

1. Background

The cultural geometry of work and gender in American English discourse has been mapped with diachronic word embeddings (Garg et al. 2018; Kozlowski et al. 2019; Jiang 2025), generally reporting a feminization-devaluation co-movement. As a concept becomes culturally associated with women, its prestige declines. Yet two limitations are pervasive in this literature. First, the analytical vocabulary is hand-curated, making the headline findings sensitive to unstated researcher choices. Second, the unit of analysis is typically pooled across the whole vocabulary, obscuring domain-level heterogeneity that sociological accounts of gendered space (Spain 1992; Massey 1994) would predict. We apply the cultural-geometry framework to physical and social spaces such as kitchens, offices, factories, churches, parks, hospitals, and prisons, using a pre-registered vocabulary protocol and a domain-stratified analytic strategy, with vocabulary emergence as a third independent line of evidence.

2. Data and vocabulary protocol

We use decade-specific aligned word2vec embeddings on Google Books American English (1900–2009, 11 decades, around 127 billion words; skip-gram with negative sampling, dim=300, min_count=100, 5 epochs, orthogonal Procrustes-aligned to a 2000s reference). Vocabulary is built through a five-stage pre-registered protocol. Stage 1 anchors on the ATUS where-category typology, the gendered-space frameworks of Spain (1992) and Massey (1994), and WordNet hyponym walks from 15 physical-space root synsets. Stage 2 recursively collects hyponym lemmas at depth four, yielding 712 raw candidates. Stage 3 applies seven pre-registered exclusion lists covering abstract locations, administrative geography, person and role nouns, polysemous tokens with dominant non-spatial senses, activity nouns, slang, and empirically identified false positives; it also adds 14 manually justified canonical tokens and 15 documented domain reassignments. Stage 4 applies a data-driven frequency filter, partitioning tokens into a main panel (present in at least 8 of 11 decades, 87 tokens) and a supplementary panel (4 to 7 of 11 decades, 78 tokens). Stage 5 merges singular and plural forms and canonicalizes spelling variants. Tokens are organized into 11 theoretical domains. Six cultural axes (prestige, potency, income, education, morality, activity) are constructed from antonym-pair centroids following Kozlowski et al. (2019); each axis measures the linguistic-discursive position of a token relative to its two anchor poles.

3. Findings

We estimate within-space two-way fixed-effects (TWFE) regressions of each of the six cultural axes on the gender axis, with space and decade fixed effects and standard errors clustered by space.

Finding 1. At the pooled main-panel level (n = 87 spaces, N = 895 observations), only the potency axis shows a marginally significant effect (beta = -0.071, p = 0.068). The prestige axis pooled effect is small and not significant (beta = -0.054, p = 0.28); the income, education, morality, and activity axes are all close to zero (|beta| < 0.013, p > 0.75). At the population level, the headline feminization-devaluation co-movement on the prestige axis does not survive systematic vocabulary expansion.

Finding 2. The pooled prestige null is the average of opposing domain-level effects. Workplace spaces (mill, factory, office, laboratory, store, mine, dairy, farm, bank, arsenal, workshop) follow a feminization-devaluation pattern (beta = -0.216, p = 0.078). Healthcare spaces (hospital, surgery, ward) follow the reverse pattern (beta = +0.265, p < 0.001), with feminization predicting prestige rise rather than decline, consistent with the historical literature on 20th-century hospital professionalization. The marginal pooled potency effect, by contrast, reflects a uniform cross-domain pattern. Five of ten domains show significant gender-potency relationships; four are negative (Vehicle beta = -0.349, p = 0.005; Outdoor home beta = -0.268, p = 0.083; Civic beta = -0.232, p = 0.005; Healthcare beta = -0.126, p < 0.001), and only Workplace is positive (beta = +0.239, p = 0.038). Potency is the most systematic cross-domain trace of the gender-axis dynamic. On a third axis constructed from rich-poor anchor word pairs (Kozlowski et al. 2019), Leisure spaces (beta = +0.195, p = 0.027) and Civic spaces (beta = +0.250, p = 0.022) feminize alongside stronger linguistic association with rich-pole anchor words.

Finding 3. Of the 78 supplementary tokens, 44 first pass the frequency threshold at or after the 1950s. These late-emergent tokens concentrate in private interior categories (kitchen 1960, bedroom 1960, classroom 1950, bathroom 1970), institutional interiors (clinic, asylum), commercial leisure (restaurant 1970, motel, salon), and specialized public spaces (stadium 1960, courthouse). Of the 44, 40 (91 percent) feminize after entry and 34 (77 percent) lose prestige after entry; 30 (68 percent) do both. Domain stratification mirrors Finding 2. Healthcare late-arrivals (clinic, asylum) are the only sub-cohort that does not lose prestige post-entry (mean delta prestige = +0.016 per decade), independently confirming the main-panel Healthcare

reversal.

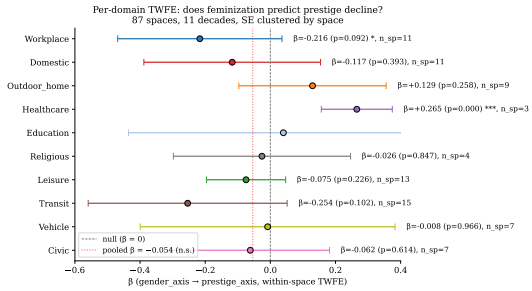

Figure 1: Per-domain TWFE coefficient of feminization on the prestige axis with 95 percent confidence intervals clustered by space. The pooled estimate (beta = -0.054, n.s.) is shown as the red dashed reference line. Workplace lies below zero (beta = -0.216, p = 0.078); Healthcare lies above zero (beta = +0.265, p < 0.001). The two domains move in opposite directions and cancel in the pooled estimate.

References

1. Garg, N., Schiebinger, L., Jurafsky, D., & Zou, J. (2018). Word embeddings quantify 100 years of gender and ethnic stereotypes. *PNAS*, 115(16), E3635–E3644.
2. Jiang, W. (2025). The cultural devaluation of feminized work. *American Sociological Review*, 90(5), 755–787.
3. Kozlowski, A. C., Taddy, M., & Evans, J. A. (2019). The geometry of culture. *American Sociological Review*, 84(5), 905–949.
4. Massey, D. (1994). *Space, Place, and Gender.* University of Minnesota Press.
5. Michel, J.-B., et al. (2011). Quantitative analysis of culture using millions of digitized books. *Science*, 331(6014), 176–182.
6. Mikolov, T., Sutskever, I., Chen, K., Corrado, G., & Dean, J. (2013). Distributed representations of words and phrases and their compositionality. *NeurIPS*, 26.
7. Pechenick, E. A., Danforth, C. M., & Dodds, P. S. (2015). Characterizing the Google Books corpus. *PLOS ONE*, 10(10).
8. Spain, D. (1992). *Gendered Spaces.* University of North Carolina Press.