From pixels to processing:
Visual and semantic predictors of word reading times

Differences in word reading times are often described in terms of experimental categories, although such categories may conflate multiple properties of lexical form and meaning. Computational models of the mental lexicon offer an alternative by deriving continuous measures of lexical organization that may provide a more direct account of processing behaviour. Here, we investigate whether such measures can explain apparent categorical differences in the reading of German gender-inclusive person nouns.

We use Linear Discriminative Learning (LDL; Baayen et al., 2019), which models lexical comprehension as learned mappings between form and semantic representations. Conventional LDL implementations typically represent written forms through letter trigrams. For gender-inclusive forms such as *Florist*in* 'florist of any gender', however, this representation reduces visually distinctive symbols such as the asterisk, colon, and underscore to discrete orthographic units and therefore abstracts away from their visual properties. We instead represent words as continuous vectors of pixel intensities derived from rendered word forms. This representation preserves visual differences while avoiding predefined graphematic segmentation. Pixel-based models performed at least as well as conventional trigram-based models and, on several measures, outperformed them.

We evaluated the model against behavioural data from a self-paced reading experiment by Haan (2025), in which 119 participants read sentences containing 24 person-denoting nouns in four conditions: generic masculine, specific masculine, specific feminine, and gender-star forms. The original analysis revealed a categorical condition effect: specific masculines were read faster than specific feminines and gender-star forms, whereas generic masculines did not differ significantly from gender-star forms.

Rather than treating these conditions as explanatory categories, we derived continuous measures of lexical processing from the pixel-based LDL model. Two measures significantly predicted target-word reading times (see Figure 1). Comprehension quality, quantifying how accurately the model reconstructs a word's semantic representation from its form, was negatively associated with reading time: words with better-reconstructed meanings were read faster. Form-neighbourhood density, quantifying the similarity of a word's predicted form to its nearest formal neighbours, likewise facilitated processing: words situated in denser form neighbourhoods were read faster. Importantly, once these lexical measures were included, the categorical condition predictor no longer explained additional variance in reading times.

These findings suggest that apparent category effects in reading can emerge from more basic, continuously varying properties of lexical form and meaning. In the present case, the processing differences between masculine, feminine, and gender-star forms can be decomposed into differences in semantic discriminability and formal neighbourhood structure rather than requiring the categories themselves as explanatory constructs. More broadly, the results show how visually grounded form representations can connect perceptual properties of written words, lexical organization, and behavioural processing within a single computational framework. Pixel-based representations therefore provide not only a less pre-structured alternative to symbolic n-grams, but also a means of moving computational models of visual word recognition closer to the information available during reading.

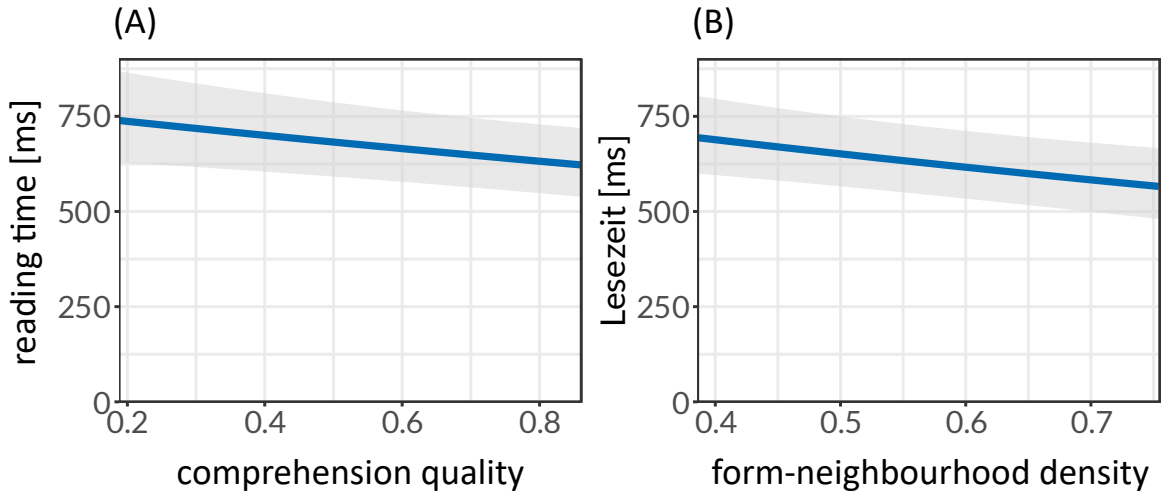


Figure 1. Predicted reading times as a function of LDL-derived comprehension quality (A) and form-neighbourhood density (B). Higher values of both measures are associated with shorter reading times.

References

Baayen, R. H., Chuang, Y.-Y., Shafaei-Bajestan, E., & Blevins, J. P. (2019). The discriminative lexicon: A unified computational model for the lexicon and lexical processing in comprehension and production grounded not in (de)composition but in linear discriminative learning. *Complexity*, 2019, 4895891. https://doi.org/10.1155/2019/4895891

Haan, Ann-Sophie. 2025. Gender-fair language processing: A self-paced reading study on the German gender star. MA thesis, Heinrich Heine University Düsseldorf.