

Understanding or Memorizing? A Case Study of German Definite Articles in Language Models

Jonathan Drechsel and Erisa Bytyqi and Steffen Herbold

Faculty of Computer Science and Mathematics

University of Passau

Correspondence: jonathan.drechsel@uni-passau.de

Abstract

Language models perform well on grammatical agreement, but it is unclear whether this reflects rule-based generalization or memorization. We study this question for German definite singular articles, whose forms depend on gender and case. Using GRADIEND, a gradient-based interpretability method, we learn parameter update directions for gender-case specific article transitions. We find that updates learned for a specific gender-case article transition frequently affect unrelated gender-case settings, with substantial overlap among the most affected neurons across settings. These results argue against a strictly rule-based encoding of German definite articles, indicating that models at least partly rely on memorized associations rather than abstract grammatical rules.

1 Introduction

Modern Language Models (LMs; Vaswani et al. 2017) achieve a near-perfect accuracy on many grammatical phenomena, yet it remains unclear *how* this competence is realized internally (Rogers et al., 2020; Belinkov and Glass, 2019; Lindsey et al., 2025). Do LMs encode abstract grammatical rules, or do they rely on surface-level memorization of frequent token-context associations? This question is particularly interesting for morphologically rich languages such as German, where grammatical gender, case, and number jointly determine surface forms (Seeker and Kuhn, 2013). Crucially, **German definite singular articles** are syncretic: the same article can appear across multiple genders and cases (e.g., *der* appears as nominative masculine and as dative/genitive feminine; see Table 1). This ambiguity lets us test whether article behavior reflects rule-based generalization or context-specific memorization, framed through two hypotheses.

H1 Memorization hypothesis: LMs memorize surface-level grammatical associations without utilizing the underlying rules.

	Nom.	Acc.	Dat.	Gen.
Male	der	den	dem	des
Neutral	das	das	dem	des
Female	die	die	der	der

Table 1: German definite singular articles.

H2 Rule-encoding hypothesis: LMs generate text based on internally represented abstract grammatical rules.

To investigate these hypotheses, we apply GRADIEND (Drechsel and Herbold, 2026), a simple encoder-decoder gradient-based interpretability method for feature learning based on parameter update directions for controlled substitutions (here, gender-case specific article swaps like *die* → *der* or *der* → *die*). GRADIEND learns a one-dimensional latent feature: its encoder maps gradients from opposite swap directions to different scalar values (± 1), while neutral inputs are encoded near 0. This scalar is then used to reconstruct gradients, which can be applied to rewrite the base model along the learned feature direction. By learning such features for different gender-case pairs, we analyze how grammatical information for article prediction is encoded internally and whether these transition-specific updates generalize across grammatical settings. See Figure 1 for an overview.

Our analysis examines (i) how applying a learned gradient direction affects article probabilities beyond the specific trained gender–case transition, and (ii) the overlap among the most affected model parameters across gender–case settings. We find statistically significant generalization across gender and case, as well as substantial neuron overlap between different transformations. Overall, our results argue against a strictly rule-based encoding of German definite articles (H2), indicating that in some contexts and nouns, article prediction is learned via memorized associations (H1) rather than abstract grammatical rules.

For brevity, we will use *article* to refer exclusively to *German definite singular articles*.

2 Related Work

2.1 Morphosyntactic Information in Model Representations

A large body of work asks whether transformers encode linguistic information internally. Probing studies suggest that syntactic structure is recoverable from representations, including hierarchical relations captured by structural probes (Hewitt and Manning, 2019) and a layer-wise organization resembling a classical NLP pipeline (Tenney et al., 2019). However, probe accuracy is not mechanistic evidence: the presence of a feature in model representations does not entail that it causally drives the model’s predictions (Belinkov, 2022).

Beyond English, multilingual probing shows that morphosyntactic features such as case, gender, and number are often accessible in model representations, with substantial variation across languages and phenomena (Acs et al., 2022). Recoverability depends on how directly and unambiguously a feature is realized in surface form. German case is explicitly identified as difficult because nouns are not case-inflected and case is marked on articles that jointly encode case and gender under high syncretism. In gender-marking languages, noun representations also exhibit distributional traces of grammatical gender (e.g., nouns sharing gender are closer in embedding space) (Gonen et al., 2019), but such effects doesn’t imply rule-based use.

2.2 Behavioral and Mechanistic Analyses of Grammar

Controlled minimal pairs provide fine-grained behavioral tests of grammatical sensitivity in LMs. Early studies show that models can prefer grammatical continuations over minimally perturbed alternatives, yet show systematic failures as constructions become more complex (Linzen et al., 2016; Marvin and Linzen, 2018). Large-scale benchmarks such as BLiMP reveal wide variation across phenomena (Warstadt et al., 2020), leaving open whether correct behavior reflects abstract rules or surface heuristics and memorized patterns.

To move beyond behavior, causal and mechanistic work intervenes on internal representations. Finlayson et al. (2021) show that modifying internal representations yields systematic changes in subject–verb agreement predictions, indicating that

grammatical behavior depends on specific internal states. Relatedly, Ferrando and Costa-jussà (2024) find highly similar circuit structures for subject–verb agreement across languages despite surface-level topological differences.

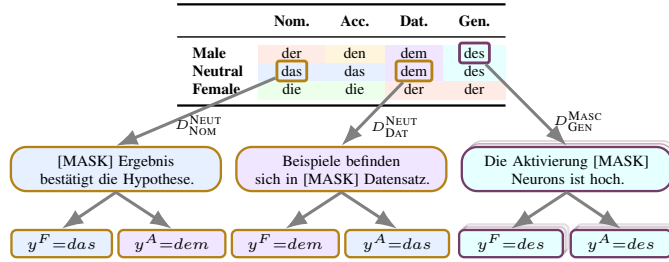
Recent sparse autoencoder (SAE) approaches (Bricken et al., 2023) further decompose activations into sparse features: Brinkmann et al. (2025) identify multilingual features corresponding to morphosyntactic concepts such as number, gender, and tense, and Jing et al. (2025) introduce *LinguaLens*, combining SAE features with counterfactual datasets and interventions to identify and manipulate mechanisms across linguistic phenomena. These results suggest that grammatical concepts can align with reusable internal feature directions. We complement this line by testing how article-transition interventions distribute across the German gender-case paradigm.

2.3 Memorization vs. Generalization

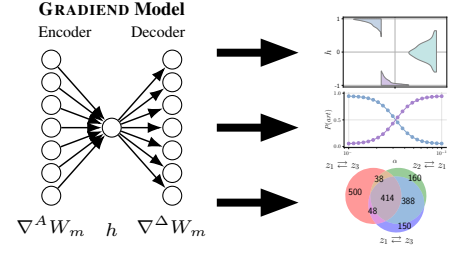
A separate line of work documents that neural LMs can *memorize* training sequences in ways that enable verbatim extraction, and that memorization increases with scale and with training-data duplication (Carlini et al., 2023). For grammar, *morphological productivity* offers a controlled test of rule-like generalization beyond frequent lexical items, e.g., Wug-style evaluations show uneven morphological generalization even for strong LMs (Weissweiler et al., 2023). Complementarily, Anh et al. (2024) find that generalization to nonce words varies systematically across languages and is predicted by morphological complexity. Together, these findings motivate our case study: high surface-level agreement can coexist with non-uniform generalization, and our gradient-based interventions probe whether German article behavior reflects reusable grammatical variables or surface-level associations.

3 Methodology

Our goal is to probe how a LM chooses German definite articles (see Table 1). We do this by asking a counterfactual question: *if the article in a given context were different, which parameters would need to change to make the model prefer the alternative article?* Concretely, we use a factual target article and a controlled counterfactual (an article swap between two gender-case cells). The resulting gradient difference isolates an update direction associated with this specific grammati-



(a) Illustration of factual (y^F) and alternative (y^A) targets for the gender-case transition (NEUT, NOM) \leftrightarrow (NEUT, DAT). Non-target cells form identity pairs (only one shown). Dataset labels (e.g., D_{NOM}^{NEUT}) denote the corresponding gender-case datasets (Section 4).



(b) GRADIEND is trained for a given transition using alternative gradients (∇^A) and the difference to factual gradients (∇^Δ). We evaluate the model via (i) encoder analysis, (ii) decoder-based probability shifts, and (iii) weight overlap across transitions.

Figure 1: Overview: Using gender-case-specific data, we learn GRADIEND models from MLM/CLM gradients of factual vs. alternative targets for two gender-case cells. GRADIEND learns to map input gradients to a scalar label h (± 1 for the two target classes; 0 for neutral inputs) and enables analyses to test our hypotheses.

cal transition. The GRADient ENcoder Decoder (GRADIEND; Drechsel and Herbold 2026) operationalizes this idea by compressing gradients into a single scalar feature and decoding it back into an update direction. This enables three analyses (see Figure 1): (i) **Encoder analysis**: whether and how the learned feature separates gradients from opposite swap directions and neutral inputs, revealing similarities across inputs; (ii) **Probability shifts**: applying decoded update directions to test which article probabilities change (and which remain stable), linking interventions to our hypotheses; and (iii) **Weight overlap**: comparing learned directions across transitions to quantify parameter reuse between conceptually different updates.

3.1 German Definite Articles as a Controlled Morphosyntactic System

German articles form a small closed-class paradigm whose surface form is determined by grammatical gender and case. Male, neutral, and female gender labels are represented by $\mathcal{G} := \{\text{MASC}, \text{NEUT}, \text{FEM}\}$, and the German cases nominative, accusative, dative, and genitive are represented by $\mathcal{C} := \{\text{NOM}, \text{ACC}, \text{DAT}, \text{GEN}\}$. We represent each gender-case combination as a *cell* $z = (g, c) \in \mathcal{G} \times \mathcal{C}$ and denote its article by $a(g, c) \in \mathcal{A}$ (defined by Table 1) with $\mathcal{A} = \{der, die, das, den, dem, des\}$. Due to *syncretism*, multiple (g, c) pairs share the same surface article (e.g., $a(\text{MASC}, \text{NOM}) = der = a(\text{FEM}, \text{DAT})$). This lets us test whether models condition article choice on abstract (g, c) variables or on surface-level token-context associations.

3.2 Article Prediction Task

We study LMs in a Masked Language Modeling (MLM; Devlin et al. 2018) setting using an article as masked target. Given a sentence, we construct an input by masking every article occurrence corresponding to a targeted gender-case cell $z = (g, c)$, while leaving the remaining context unchanged.

For each masked instance, we define two targets: (i) **Factual target** y^F , the grammatically licensed article for $z = (g, c)$ specified by the sentence context, i.e., $y^F = a(g, c)$. (ii) **Alternative target** y^A , an article specified by a predefined transition between two cells (defined in the next subsection).

These induce corresponding factual ($\nabla^F W_m$) and alternative ($\nabla^A W_m$) gradients with respect to the selected model parameters W_m . We define their difference as $\nabla^\Delta W_m := \nabla^F W_m - \nabla^A W_m$.

3.3 GRADIEND for German Gender

We train one GRADIEND model (Figure 1b) per targeted transition $T = (z_1 \leftrightarrow z_2)$ between gender-case cells $z_i = (g_i, c_i)$ differing in exactly one dimension: gender at fixed case ($g_1 \neq g_2, c_1 = c_2$) or case at fixed gender ($g_1 = g_2, c_1 \neq c_2$). For instance, the nominative gender transition $z_1 = (\text{MASC}, \text{NOM})$ and $z_2 = (\text{FEM}, \text{NOM})$ corresponds to the article transition $der \leftrightarrow die$. For masking tasks for z_1 and z_2 , we construct *swapped target pairs*: for z_1 we set $(y^F, y^A) = (a(z_1), a(z_2))$, and for z_2 we set $(y^F, y^A) = (a(z_2), a(z_1))$. These swaps induce non-zero gradient differences $\nabla^\Delta W_m$ that encode the transition direction T .

To keep the learned update specific to T , we additionally include masking tasks from all other cells $z \notin \{z_1, z_2\}$ as *identity pairs*. Under this construc-

tion, the factual and alternative gradients are identical by definition, yielding $\nabla^\Delta W_m = 0$. This explicitly enforces a *do-not-change* constraint: GRADIEND is trained to produce no update for non-targeted gender-case settings.

Figure 1a illustrates the construction. We denote GRADIEND models targeting gender transitions at fixed case c as $G_c^{g_1, g_2}$ and case transitions at fixed gender g as G_{c_1, c_2}^g , e.g., $G_{\text{NOM}}^{\text{FEM, MASC}}$ and $G_{\text{NOM, DAT}}^{\text{NEUT}}$.

3.4 GRADIEND Architecture and Training

GRADIEND learns a bottleneck encoder-decoder $f = \text{dec} \circ \text{enc}$ that maps gradient information to a single scalar feature and decodes it into a parameter-space update direction (see Figure 1b). We use a one-dimensional bottleneck $h \in [-1, 1]$:

$$h = \text{enc}(\nabla_{\text{in}} W_m) = \tanh(W_e^\top \nabla_{\text{in}} W_m + b_e),$$

$$\nabla_{\text{out}} W_m \approx \text{dec}(h) = h \cdot W_d + b_d,$$

where $W_m \in \mathbb{R}^n$ are the selected model parameters and $W_e, W_d, b_d \in \mathbb{R}^n, b_e \in \mathbb{R}$ are learned.

Departing from Drechsel and Herbold (2026), who use factual gradients as input, we set $\nabla_{\text{in}} W_m := \nabla^A W_m$ and $\nabla_{\text{out}} W_m := \nabla^\Delta W_m$. Article prediction is often highly confident, making factual gradients near-zero and GRADIEND training unstable. Alternative targets of the (g, c) , by construction, typically receive substantially lower probability, producing more informative gradients.

We train GRADIEND with the reconstruction loss

$$\mathcal{L}_{\text{GRADIEND}} = \|\text{dec}(\text{enc}(\nabla_{\text{in}} W_m)) - \nabla_{\text{out}} W_m\|_2^2,$$

encouraging h to encode the targeted transition while remaining neutral for identity pairs.

After training, GRADIEND yields an update direction for any $h^* \in \mathbb{R}$ via $\text{dec}(h^*)$. We intervene on the base model with $\widetilde{W}_m = W_m + \alpha \cdot \text{dec}(h^*)$ with *learning rate* α .

4 Data

To extract gender-case-specific article transition gradients, we construct one dataset for each cell $z = (g, c) \in \mathcal{G} \times \mathcal{C}$. Each dataset contains sentences in which only articles corresponding to z are masked. We filter German Wikipedia sentences (Wikimedia Foundation, 2022), retaining only those where spaCy (Honnibal et al., 2020) identifies a definite singular article with the desired gender and case (Figure 2). We denote the resulting datasets by D_c^g (e.g., $D_{\text{NOM}}^{\text{MASC}}$ for masculine-

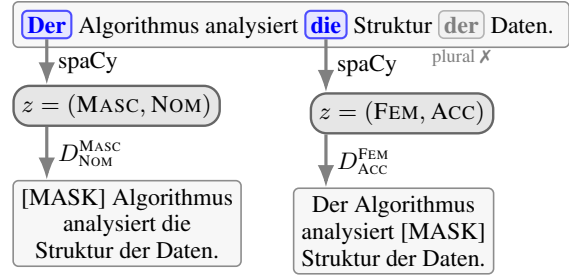


Figure 2: Data generation: spaCy determines gender and case of articles to determine the target dataset.

nominative). Sizes range from 19K to 61K (Table 6), and each dataset is split into train (80%), validation (10%), and test (10%) subsets.

To probe behavior without gender/case cues, we construct D_{NEUTRAL} , a dataset with minimized gender and case cues. It is derived from the Wortschatz Leipzig German news corpus (Goldhahn et al., 2012; Leipzig Corpora Collection, 2024) and filtered to exclude sentences containing determiners, definite or indefinite articles, or third-person pronouns. The resulting dataset serves as a gender-case-independent reference in our analyses.

Full generation details are in Appendix A.

5 Experiments

We evaluate whether German definite articles in LMs are memorized from context (H1) or determined via abstract grammatical representations (H2). We analyze GRADIEND models from three complementary angles: (i) how encoded values h distribute, (ii) how applying the decoded update affects article probabilities across gender-case cells, and (iii) how similar the learned update directions are in parameter space via Top- k weight overlap.

5.1 Experimental Setup

Models. We study four German models (GermanBERT, GBERT, ModernGBERT, GermanGPT-2) and two multilingual models (EuroBERT, LLaMA), covering encoder-only and decoder-only transformers. We include ModernGBERT (1B parameters) as an intermediate-size model between the smaller German models (109M–336M) and LLaMA (3.2B). Table 13 summarizes architectures and sizes.

Targeted transitions. Across the German article paradigm, transitions cluster into: (i) *two-dimensional* groups (gender- and case-based transitions within the same article pair), (ii) *one-dimensional* groups (multiple transitions along a single dimension), and (iii) *singleton* groups (a single

Article Pair	Datasets		GRADIEND Variants	
	Left Article	Right Article	Gender Transition	Case Transition
Two-dimensional transitions (gender <i>and</i> case vary)				
<i>der</i> ↔ <i>die</i>	$D_{\text{NOM}}^{\text{MASC}}, D_{\text{DAT}}^{\text{FEM}}, D_{\text{GEN}}^{\text{FEM}}$	$D_{\text{NOM}}^{\text{FEM}}, D_{\text{ACC}}^{\text{FEM}}$	$G_{\text{NOM}}^{\text{FEM,MASC}}$	$G_{\text{NOM,DAT}}^{\text{FEM}}, G_{\text{NOM,GEN}}^{\text{FEM}}, G_{\text{ACC,DAT}}^{\text{FEM}}, G_{\text{ACC,GEN}}^{\text{FEM}}$
<i>der</i> ↔ <i>dem</i>	$D_{\text{NOM}}^{\text{MASC}}, D_{\text{DAT}}^{\text{FEM}}$	$D_{\text{DAT}}^{\text{MASC}}, D_{\text{DAT}}^{\text{NEUT}}$	$G_{\text{DAT}}^{\text{FEM,MASC}}, G_{\text{DAT}}^{\text{FEM,NEUT}}$	$G_{\text{NOM,DAT}}^{\text{MASC}}$
<i>der</i> ↔ <i>des</i>	$D_{\text{NOM}}^{\text{MASC}}, D_{\text{GEN}}^{\text{FEM}}$	$D_{\text{GEN}}^{\text{MASC}}, D_{\text{GEN}}^{\text{NEUT}}$	$G_{\text{GEN}}^{\text{FEM,MASC}}, G_{\text{GEN}}^{\text{FEM,NEUT}}$	$G_{\text{NOM,GEN}}^{\text{MASC}}$
One-dimensional transitions (only gender <i>or</i> only case varies)				
<i>das</i> ↔ <i>die</i>	$D_{\text{NOM}}^{\text{FEM}}, D_{\text{ACC}}^{\text{FEM}}$	$D_{\text{NOM}}^{\text{NEUT}}, D_{\text{ACC}}^{\text{NEUT}}$	$G_{\text{NOM}}^{\text{FEM,NEUT}}, G_{\text{ACC}}^{\text{FEM,NEUT}}$	
<i>das</i> ↔ <i>dem</i>	$D_{\text{NOM}}^{\text{NEUT}}, D_{\text{ACC}}^{\text{NEUT}}$	$D_{\text{DAT}}^{\text{NEUT}}$		$G_{\text{NOM,DAT}}^{\text{NEUT}}, G_{\text{ACC,DAT}}^{\text{NEUT}}$
<i>das</i> ↔ <i>des</i>	$D_{\text{NOM}}^{\text{NEUT}}, D_{\text{ACC}}^{\text{NEUT}}$	$D_{\text{GEN}}^{\text{NEUT}}$		$G_{\text{NOM,GEN}}^{\text{NEUT}}, G_{\text{ACC,GEN}}^{\text{NEUT}}$
<i>dem</i> ↔ <i>des</i>	$D_{\text{DAT}}^{\text{MASC}}, D_{\text{DAT}}^{\text{NEUT}}$	$D_{\text{GEN}}^{\text{MASC}}, D_{\text{GEN}}^{\text{NEUT}}$		$G_{\text{GEN,DAT}}^{\text{MASC}}, G_{\text{GEN,DAT}}^{\text{NEUT}}$

Table 2: Targeted bidirectional article transitions and GRADIEND variants. Listed are all trained transitions grouped by their structural diversity (two- vs. one-dimensional), together with the corresponding datasets and model variants.

transition). We focus on two- and one-dimensional groups, which enable within-group comparisons, and list all these transitions in Table 2. For each of these transitions $T = (z_1 \leftrightarrow z_2)$, we train one GRADIEND model, inducing two *directed* transitions $z_1 \rightarrow z_2$ and $z_2 \rightarrow z_1$.

Training. We train GRADIEND models as described in Section 3, using swapped targets for the two cells defining T and identity pairs for all remaining cells, with the gender-case datasets from Section 4. For consistent visualization across base models, we normalize the sign of the encoded value so the same targeted directional article transition has a consistent polarity (positive vs. negative h). Training details are provided in Appendix C.

Evaluation. Unless stated otherwise, we evaluate on the test splits of the corresponding datasets.

Decoder-only models. German articles depend on the noun’s gender, so left-to-right context is insufficient. Hence, we add a MLM-style article classifier for bidirectional conditioning (Appendix B).

5.2 Feature Encoding Analysis

We analyze how GRADIEND maps gradient inputs to the scalar bottleneck value h . Figure 3 shows the encoded-value distributions for a representative GRADIEND variant, $G_{\text{NOM}}^{\text{FEM,MASC}}$, across all base models (other variants in Appendix D). Table 3 complements this view by reporting correlations between h and our expected labels, assigning ± 1 to the two directed transition tasks and 0 to identity pairs (neutral updates). All models reach correlations of at least 50% across all GRADIENDs, while German encoder-only models often exceed 90%.

Stable orientation on the targeted transition. Across models, the two targeted directed transitions (blue in Figure 3) map to opposite signs, consistently




Figure 3: Encoded value distribution of $G_{\text{NOM}}^{\text{FEM,MASC}}$ (other GRADIENDs in Figures 11-28).

tently separating the directional article transitions.

Other realizations of the same article transition. Beyond the trained pair (z_1, z_2) , we evaluate the encoder on other gender-case pairs $(\tilde{z}_1, \tilde{z}_2)$ that realize the same article transitions $a(\tilde{z}_1) \rightarrow a(\tilde{z}_2)$ and vice versa (green). Encoder-only models often assign these non-target transitions the same signed encoding as the trained transition, suggesting that gradient directions for different realizations of an article transition are closely aligned. Decoder-only models show less stable behavior, with encodings frequently clustering around zero rather than the extremes ± 1 , probably due to the custom MLM-style prediction head used during GRADIEND training.

Identity pairs. By construction, identity-pair gradients (red/orange/purple) map to $h \approx 0$, which is clearly observed for German encoder-only mod-

	<i>der</i> ↔ <i>die</i>					<i>der</i> ↔ <i>dem</i>			<i>der</i> ↔ <i>des</i>			<i>das</i> ↔ <i>die</i>		<i>das</i> ↔ <i>dem</i>		<i>das</i> ↔ <i>des</i>		<i>dem</i> ↔ <i>des</i>	
	$G_{\text{NOM}}^{\text{FEM,MASC}}$	$G_{\text{NOM,DAT}}^{\text{FEM}}$	$G_{\text{NOM,GEN}}^{\text{FEM}}$	$G_{\text{ACC,DAT}}^{\text{FEM}}$	$G_{\text{ACC,GEN}}^{\text{FEM}}$	$G_{\text{DAT}}^{\text{FEM,MASC}}$	$G_{\text{DAT}}^{\text{FEM,NEUT}}$	$G_{\text{NOM,DAT}}^{\text{MASC}}$	$G_{\text{GEN}}^{\text{FEM,MASC}}$	$G_{\text{GEN}}^{\text{FEM,NEUT}}$	$G_{\text{NOM,GEN}}^{\text{MASC}}$	$G_{\text{ACC}}^{\text{FEM,NEUT}}$	$G_{\text{NOM}}^{\text{FEM,NEUT}}$	$G_{\text{NOM,DAT}}^{\text{NEUT}}$	$G_{\text{ACC,DAT}}^{\text{NEUT}}$	$G_{\text{NOM,GEN}}^{\text{NEUT}}$	$G_{\text{ACC,GEN}}^{\text{NEUT}}$	$G_{\text{GEN,DAT}}^{\text{MASC}}$	$G_{\text{GEN,DAT}}^{\text{NEUT}}$
GermanBERT	96.3	95.7	97.1	95.3	91.9	96.4	95.8	95.7	96.5	95.7	96.6	93.5	97.8	96.5	90.7	96.6	94.1	98.4	97.0
GBERT	98.4	97.8	96.3	96.1	96.2	98.6	98.4	97.8	97.9	97.7	95.6	96.3	98.6	96.6	95.3	97.1	97.3	98.2	97.4
ModernGBERT	94.1	93.9	93.3	84.3	82.7	90.2	88.7	93.9	87.2	85.5	92.2	88.7	95.0	93.5	84.8	91.2	81.3	94.6	90.9
EuroBERT	61.9	61.3	59.4	52.2	50.6	61.2	59.8	61.3	50.7	50.6	64.7	64.7	72.9	61.7	52.3	65.7	55.1	66.2	53.3
GermanGPT-2	56.6	64.3	58.8	57.8	51.0	69.1	59.5	64.3	61.7	58.0	67.3	58.2	71.1	63.1	57.1	62.5	67.7	65.6	57.1
LLaMA	59.2	62.1	58.5	55.8	52.5	61.3	55.2	62.1	58.9	51.1	57.6	52.9	67.0	57.7	50.5	62.9	54.6	63.9	55.8

Table 3: Pearson correlation of encoded values, scaled by 100.

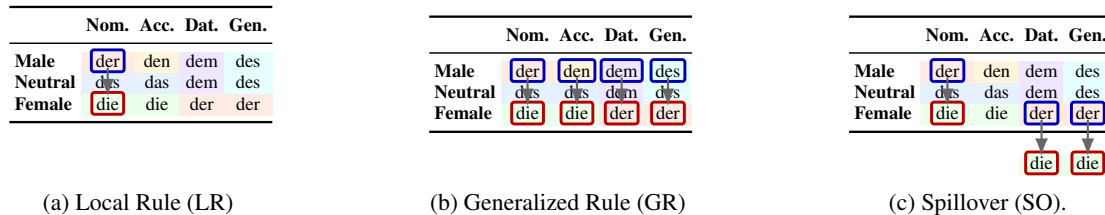


Figure 4: Patterns of generalizations, exemplified using $G_{\text{NOM}}^{\text{FEM,MASC}}$ (*der* → *die*).

els. EuroBERT exhibits larger deviations: while most identity pairs remain centered near zero, gradients involving articles from the trained GRADIEND variant shift toward the same signed encoding as the targeted transition, aligning with the sign of the target article. A plausible explanation is that multilingual representations encode German article distinctions less sharply, so gradients for factual predictions may still share task-relevant information in the article prediction task. Decoder-only models show a similar pattern and additionally spread identity pairs involving other articles across much of $[-1, 1]$, likely reflecting variance introduced by our lightweight article classification head.

Neutral control. Finally, gradients on D_{NEUTRAL} (where no articles are masked by construction) map consistently close to zero across models and configurations (yellow), providing a sanity check.

GRADIEND learns a meaningful scalar representation that robustly separates transition directions and often generalizes across gender-case pairs of the same surface article transition.

5.3 Intervention Effects on Articles

Next, we evaluate how GRADIEND updates affect article probabilities and how these effects distribute across gender-case cells. Examining where probability changes occur is central to our analysis, since different internal mechanisms imply different patterns of generalization. A grammar-tracking mechanism should yield either (i) *local rule-based* (LR)

effects restricted to the trained cell, or (ii) *generalized rule-based* (GR) effects that propagate systematically to grammatically related cells (e.g., along gender while preserving case). In contrast, surface-level behavior can yield *spillover* (SO), where the same surface transition (e.g., *der* → *die*) appears in grammatically unrelated cells that share the same source article. Figure 4 illustrates these patterns.

Selected article transitions. To focus on the most diagnostic settings, we restrict this analysis to article groups containing gender and case transitions, enabling evaluation of a trained GRADIEND along the other dimension (Table 2). Such evaluations are possible within each two-dimensional transition group. For example, we assess the impact of $G_{\text{NOM}}^{\text{FEM,MASC}}$ *der* → *die* on $D_{\text{DAT}}^{\text{FEM}}$ and $D_{\text{GEN}}^{\text{FEM}}$, since these datasets share the source article *der*.

Intervention strength and α selection. Large updates can change predictions by degrading language modeling (Drechsel and Herbold, 2026). Figure 5 shows this trade-off for GermanBERT under $G_{\text{NOM,DAT}}^{\text{FEM}}$ *der* → *die*: as α increases, $\mathbb{P}(\textit{die})$ rises (until a certain point) while a Language Modeling Score (LMS) measured on D_{NEUTRAL} drops. We therefore analyze probability shifts only under an explicit LMS-preservation constraint. Concretely, we apply scaled decoder updates $\widetilde{W}_m = W_m + \alpha \cdot \textit{dec}(h^*)$, where $h^* = \pm 1$ selects the transition direction, and evaluate a grid of $\alpha > 0$ values. We retain only candidates that preserve at least 99% of the base-model LMS on D_{NEUTRAL} (masked-token accuracy for encoder-only models, perplexity

Model	Art. Trans.	α	$D_{\text{NOM}}^{\text{MASC}}(der)$			$D_{\text{GEN}}^{\text{FEM}}(der)$			$D_{\text{DAT}}^{\text{FEM}}(der)$			D_{NEUTRAL}			SuperGLEBer
			$\Delta\mathbb{P}$	d	Sig.	$\Delta\mathbb{P}$	d	Sig.	$\Delta\mathbb{P}$	d	Sig.	$\Delta\mathbb{P}$	d	Sig.	
GermanBERT	—	0.0	—	—	—	—	—	—	—	—	—	—	—	—	70.7 ± 0.4
+ $G_{\text{NOM}}^{\text{FEM,MASC}}$	<i>der</i> → <i>die</i>	0.01	0.04	0.30	***	0.00	0.12	***	0.01	0.16	***	0.05	0.01	n.s.	70.1 ± 0.4
+ $G_{\text{DAT}}^{\text{FEM}}$	<i>der</i> → <i>die</i>	0.1	0.05	0.25	***	0.04	0.16	***	0.33	0.24	***	0.09	0.02	n.s.	70.2 ± 0.4
+ $G_{\text{GEN}}^{\text{FEM}}$	<i>der</i> → <i>die</i>	0.5	0.66	0.32	***	1.85	0.47	***	1.40	0.37	***	0.18	0.03	**	70.2 ± 0.4
+ $G_{\text{DAT}}^{\text{MASC}}$	<i>der</i> → <i>dem</i>	0.05	0.04	0.16	***	0.00	0.08	***	0.02	0.13	***	0.02	0.01	n.s.	70.1 ± 0.4
+ $G_{\text{GEN}}^{\text{FEM,NEUT}}$	<i>der</i> → <i>dem</i>	0.01	0.00	0.11	***	0.00	0.08	***	0.01	0.17	***	0.03	0.02	n.s.	70.1 ± 0.4
+ $G_{\text{DAT}}^{\text{MASC}}$	<i>der</i> → <i>des</i>	0.2	0.14	0.15	***	0.28	0.35	***	0.02	0.04	***	0.01	0.01	n.s.	70.2 ± 0.4
+ $G_{\text{GEN}}^{\text{FEM,NEUT}}$	<i>der</i> → <i>des</i>	0.5	0.25	0.12	***	3.70	0.55	***	0.08	0.07	***	0.07	0.05	***	70.2 ± 0.4

Table 4: GRADIEND-modified GermanBERT models (others in Table 15): $\Delta\mathbb{P}$ of target article (scaled by 100), Cohen’s d , and significance as *** ($p < .001$, ** $p < .01$, * $p < .05$; n.s. otherwise). Bold marks corresponding GRADIEND datasets. SuperGLEBer score (scaled by 100) use bootstrapped 95% confidence intervals ($n = 1000$).




Figure 5: $G_{\text{NOM,DAT}}^{\text{FEM}}$ applied to GermanBERT for the *der* → *die* transition: mean article probabilities and LMS across learning rates α . The candidate range ($\alpha > 0$ and before the 99% LMS drop) is shaded gray. Highlighted LMS points mark the base model (left) and α^* (maximizing $\mathbb{P}(der)$ on $D_{\text{DAT}}^{\text{FEM}}$ in gray area).

for decoder-only models), and choose α^* as the candidate that maximizes the mean probability of the target article on the corresponding target-article dataset (e.g., $D_{\text{DAT}}^{\text{FEM}}$ for $G_{\text{NOM,DAT}}^{\text{FEM}}$ *der* → *die*). The candidate range and α^* are highlighted in Figure 5. Details are in Appendix E.2.

Probability evaluation. For each gender-case dataset, we compute the mean article-probability change $\Delta\mathbb{P}(art)$ of the GRADIEND-modified model relative to the base (positive indicates an increase). We report Cohen’s d with pooled variance (Cohen, 1988) and test significance with a permutation test (Good, 2005). Table 4 shows representative results for GermanBERT (other models in Table 15).

Effects occur before broad degradation.

Across transitions, GRADIEND-updates induce significant shifts on article datasets while changes on D_{NEUTRAL} are mostly non-significant. Because α^* is chosen conservatively, $\Delta\mathbb{P}$ is typically below 1%, but effect sizes and significance indicate consistent

directional shifts. Effects are usually strongest on the trained cell, yet remain substantially larger on other article datasets sharing the same source article than on D_{NEUTRAL} , suggesting the changes are not due to broad degradation. This is further supported by mostly unchanged SuperGLEBer scores, a German NLP benchmark consisting of 29 tasks (Pfister and Hotho, 2024).

Effects on all cells. Figure 6 visualizes $\Delta\mathbb{P}$ over the full gender-case grid for the $G_{\text{NOM}}^{\text{FEM,MASC}}$ (*der* → *die*) across all base models as heatmap. We additionally overlay the three patterns of generalization from Figure 4. The heatmap partially matches GR, but with deviations: some GR-predicted cells are neutral or opposite-signed (e.g., GermanBERT $D_{\text{DAT}}^{\text{FEM}}$ for *der*), and several effects appear in cells that are not predicted by a clean grammar-preserving rule. Notably, the two clearest GR contradictions in Figure 6, $\mathbb{P}(der)$ on $D_{\text{DAT}}^{\text{FEM}}$ and $D_{\text{GEN}}^{\text{FEM}}$, align with SO, which explains most of its predicted cells in terms of probability-change direction. Decoder-only models show more deviations, probably due to the custom small MLM head. LLaMA is the only model not showing the $\mathbb{P}(der)$ increase on $D_{\text{DAT}}^{\text{FEM}}/D_{\text{GEN}}^{\text{FEM}}$, possibly indicating a trend of less memorization in larger models. Across models, cells sharing a surface article with the same gender *or* case (e.g., $D_{\text{DAT}}^{\text{MASC}}/D_{\text{DAT}}^{\text{NEUT}}$ and $D_{\text{GEN}}^{\text{MASC}}/D_{\text{GEN}}^{\text{NEUT}}$) often behave similarly, indicating transitive spillover. For example, $G_{\text{NOM}}^{\text{FEM,MASC}}$ *die* → *der* increases $\mathbb{P}(des)$ on $D_{\text{GEN}}^{\text{MASC}}$ (GR consistent), and concurrently on $D_{\text{GEN}}^{\text{NEUT}}$.

Probability-shift patterns are neither fully rule-based nor consistent with unrestricted spillover, but instead suggest a mixture of structured generalization and surface-article-linked effects.

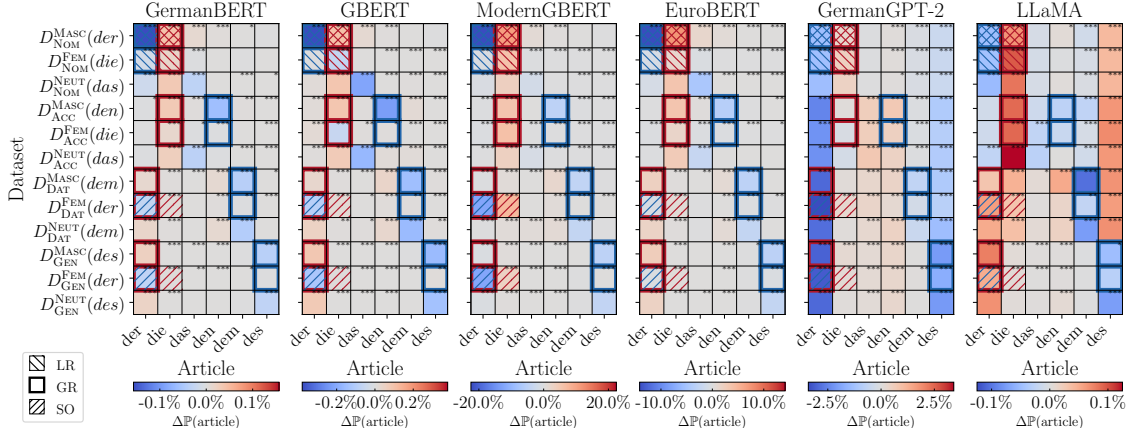


Figure 6: Mean probability change of articles between GRADIEND-modified and base model for $G_{\text{NOM}}^{\text{FEM,MASC}} \text{der} \rightarrow \text{die}$ (others in Figures 29-34). Stars mark statistical significance after Benjamini-Hochberg FDR correction (Benjamini and Hochberg, 1995) applied per model. Marked cells are expectations for LR, GR, and SO (Figure 4).




Figure 7: Top-1,000 weight overlaps across different GRADIENDs for GermanBERT (other models in Figure 35).

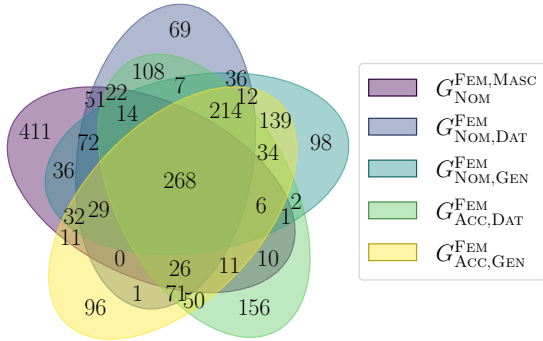


Figure 8: Top-1,000 weight overlaps for GermanBERT $\text{der} \leftrightarrow \text{die}$ article group (other models in Figure 36).




Figure 9: Top-1,000 weight overlaps for the GermanBERT control group (other models in Figure 37).

5.4 Overlap of Most Affected Parameters

Finally, we compare GRADIEND models directly in the parameter space. We define parameter importance by the absolute value of decoder weights W_d and extract the Top- k weights for $k = 1000$.

Overlap within article groups. Figures 7 and 8 show article group Venn diagrams for GermanBERT (additional models are reported in Appendix F). Across base models, Top- k weights overlap substantially across GRADIENDs within an article group. For instance, overlap in the $\text{der} \leftrightarrow \text{die}$ group remains high even between variants trained on different case/gender axes (e.g., $G_{\text{NOM}}^{\text{FEM,MASC}}$ vs. $G_{\text{NOM,DAT}}^{\text{FEM}}$ and $G_{\text{NOM,GEN}}^{\text{FEM}}$), suggesting a shared subset of weights despite differing grammatical contexts.

Control group. To test whether overlap is expected in general, we analyze a control group of four GRADIEND variants spanning ACC/DAT and FEM/NEUT whose cells realize disjoint surface articles. Figure 9 shows the Top- k overlap for GermanBERT, which is smaller than in the article groups, indicating that the high overlap in Figures 7 and 8 is not a generic artifact.

Quantifying overlap. Table 5 quantifies these observations. For each article group (including the control), we report the maximum pairwise Top- k overlap, $\max_{A,B} \frac{|A \cap B|}{k}$, where A and B are the Top- k sets of two variants. Groups with at least two variants realizing the same surface article pair show consistently high overlap ($> 75\%$ on average), whereas the control group is much lower

Model	<i>der</i> ↔ <i>die</i>	<i>der</i> ↔ <i>dem</i>	<i>der</i> ↔ <i>des</i>	<i>das</i> ↔ <i>die</i>	<i>das</i> ↔ <i>das</i>	<i>das</i> ↔ <i>des</i>	<i>dem</i> ↔ <i>des</i>	Control
GermanBERT	73.4	92.1	80.2	63.5	56.3	73.5	78.5	27.4
GBERT	77.4	86.4	76.8	64.2	81.1	61.3	69.5	27.5
ModernGBERT	87.2	89.7	70.9	80.7	70.0	68.5	66.2	27.5
EuroBERT	83.7	95.3	88.1	71.6	78.3	78.0	85.4	51.2
GermanGPT-2	90.8	93.9	86.5	85.4	86.8	83.7	95.7	51.7
LLaMA	91.7	82.1	91.5	88.1	83.6	85.3	87.2	48.2
Mean	84.0	89.9	82.3	75.6	76.0	75.1	80.4	38.9

Table 5: Maximum pairwise Top-1,000 weight overlap (scaled by 100) for article groups including the control group.

(mean 38.9%). This suggests that gender-case transitions of the same articles rely on a shared subset of parameters rather than disjoint, transition-specific mechanisms.

Top- k weights overlap strongly within article groups and far less in the disjoint-article control groups, suggesting shared parameters for the same article transitions.

5.5 Discussion

Taken together, our analyses show that German definite article behavior is not fully explained by a uniformly rule-based mechanism, providing evidence against **H2**. First, the encoding analysis shows that the learned bottleneck value h reliably separates the two swapped training cells, yet often assigns similar encodings to gradients from other gender-case transitions with the same articles, suggesting limited cell-specific disentanglement. Second, interventions shift article probabilities beyond the trained cell and only partially follow clean gender-/case-preserving generalization, while exhibiting spillover. We also observe a tentative size trend: larger models (LLaMA) show less spillover (as also suggested by the encoded value distributions, e.g., Figure 3). Third, the Top- k overlap analysis reveals substantial intersections of the most affected weights across variants within an article group, with smaller overlaps in a control group with disjoint surface articles, indicating that multiple transitions rely on a shared parameter subspace.

Our results provide evidence that a purely rule-based account (**H2**) is unlikely to be sufficient, and suggest that the memorization hypothesis (**H1**) holds at least to some extent.

6 Conclusion

We studied whether German definite singular articles in LMs reflect abstract rule encoding (**H2**) or surface-level memorization (**H1**). Using GRADIEND across multiple gender-case cells, we find that article transition updates shift article probabilities significantly beyond the trained cell under a LMS constraint and only partially follow rule-based generalization. The results are not consistent with a purely rule-based encoding of the grammar and suggest that memorization-like mechanisms are important. This suggests that while LMs can be used reliably to assess if text is grammatical and to produce grammatical text, they should be used with care when using them to analyze grammatical systems, since rules might not be encoded as expected.

We release our code and datasets: <https://github.com/aieng-lab/gradiend-german-articles>.

7 Limitations

Our study has several limitations that constrain the scope of the conclusions.

First, we focus exclusively on **German definite singular articles**, a small and highly regular closed-class system. While this makes the analysis controlled and interpretable, the findings may not transfer to other morphosyntactic phenomena (e.g., adjective agreement, verb inflection, or freer word order) or to other languages, where grammatical cues are distributed differently. However, the lack of a strict, rule-based encoding of such a regular, closed-class system indicates that more complex systems are also not learned through rules, but rather memorized.

Second, our conclusions rely on **gradient-based interventions** using GRADIEND. Although the applied update is dense and affects (in principle) all parameters, it is restricted to a single update direction scaled by a scalar α . Thus, our interventions primarily reveal mechanisms that can be expressed as a coherent, scalable update direction. More distributed or highly context-specific rule implementations may not be captured by this probe.

Third, the measured **intervention effects are small by design**. Because we select α^* under a strict LMS-preservation criterion, mean probability shifts are typically below 1%. While effect sizes and significance indicate consistent directional changes, small magnitudes make it harder

to judge downstream behavioral impact and may understate the extent of generalization that would appear under less conservative constraints.

Fourth, we evaluate only a **limited model scale range** (up to $\sim 3\text{B}$ parameters). Larger models or models trained with substantially different data mixtures or objectives may encode grammatical information differently, and scaling trends cannot be confidently concluded from our setup. Nevertheless, we highlight that all models we analyzed consistently yielded results consistent with our memorization hypothesis, indicating that this is a general pattern.

Fifth, we rely on **spaCy-based annotation** to construct gender-case datasets. While effective at scale, automatic annotation can introduce noise, especially in ambiguous or syntactically complex sentences, which may affect gradient estimates and significance patterns. To estimate labeling quality, we manually annotated 20 randomly sampled instances per gender-case dataset and obtain an overall accuracy of 82% (Appendix A.1.4). Many mis-classifications are not for other definite articles and are rather other errors that should only introduce random noise, e.g., mislabelling the relative pronoun *der* as an article. Errors that mix up definite singular articles are only commonly occurring for *des*, but there is no visible different trend when restricting to transitions and dataset intersections with near-perfect labeling, suggesting that spaCy noise is unlikely to impact our findings. This noise-tolerance of GRADIEND is further supported by earlier results, where noisy features for race and religion could also be identified by GRADIEND (Drechsel and Herbold, 2026).

Sixth, decoder-only models are evaluated using a **custom MLM-style head** to enable bidirectional conditioning. This departs from their native training objective and may influence gradient structure and intervention behavior, limiting direct comparability with encoder-only models. However, the general consistence of the results with the encoder-only models indicates a limited impact on our study of this, though we believe that this is – at least partially – responsible for different pattern for neutral results that the MLM heads are not specifically trained for.

Finally, our evaluation focuses on **controlled probability shifts** and parameter overlap rather than downstream generation behavior.

References

- Judit Acs, Endre Hamerlik, Roy Schwartz, Noah A Smith, and Andras Kornai. 2022. [Morphosyntactic probing of multilingual BERT models](#). *Natural Language Engineering*, 30.
- Dang Anh, Limor Raviv, and Lukas Galke. 2024. [Morphology matters: Probing the cross-linguistic morphological generalization abilities of large language models through a wug test](#). In *Proceedings of the Workshop on Cognitive Modeling and Computational Linguistics*, pages 177–188, Bangkok, Thailand. Association for Computational Linguistics.
- Yonatan Belinkov. 2022. [Probing classifiers: Promises, shortcomings, and advances](#). *Computational Linguistics*, 48(1):207–219.
- Yonatan Belinkov and James Glass. 2019. [Analysis methods in neural language processing: A survey](#). *Transactions of the Association for Computational Linguistics*, 7:49–72.
- Yoav Benjamini and Yosef Hochberg. 1995. [Controlling the false discovery rate: A practical and powerful approach to multiple testing](#). *Journal of the Royal Statistical Society: Series B (Methodological)*, 57(1):289–300.
- Nicolas Boizard, Hippolyte Gisserot-Boukhlef, Duarte Miguel Alves, Andre Martins, Ayoub Hamal, Caio Corro, CELINE HUDELLOT, Emmanuel Malherbe, Etienne Malaboeuf, Fanny Jourdan, Gabriel Hautreux, João Alves, Kevin El Haddad, Manuel Faysse, Maxime Peyrard, Nuno M Guerreiro, Patrick Fernandes, Ricardo Rei, and Pierre Colombo. 2025. [EuroBERT: Scaling multilingual encoders for european languages](#). In *Second Conference on Language Modeling*.
- Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Conerly, Nick Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu, Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex Tamkin, Karina Nguyen, and 6 others. 2023. [Towards monosemanticity: Decomposing language models with dictionary learning](#). *Transformer Circuits Thread*.
- Jannik Brinkmann, Chris Wendler, Christian Bartelt, and Aaron Mueller. 2025. [Large language models share representations of latent grammatical concepts across typologically diverse languages](#). In *Proceedings of the 2025 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 6131–6150, Albuquerque, New Mexico. Association for Computational Linguistics.
- Nicholas Carlini, Daphne Ippolito, Matthew Jagielski, Katherine Lee, Florian Tramèr, and Chiyuan Zhang.

2023. [Quantifying memorization across neural language models](#). In *Proceedings of the 11th International Conference on Learning Representations (ICLR 2023)*.
- Branden Chan, Stefan Schweter, and Timo Möller. 2020. [German’s next language model](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 6788–6796, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Jacob Cohen. 1988. *Statistical Power Analysis for the Behavioral Sciences*, 2 edition. Lawrence Erlbaum Associates, Hillsdale, NJ.
- A. C. Davison and D. V. Hinkley. 1997. *Bootstrap Methods and their Application*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. [BERT: pre-training of deep bidirectional transformers for language understanding](#). *CoRR*, abs/1810.04805.
- Jonathan Drechsel and Steffen Herbold. 2026. [GRA-DIEND: Feature learning within neural networks exemplified through biases](#). In *The Fourteenth International Conference on Learning Representations*.
- Javier Ferrando and Marta R. Costa-jussà. 2024. [On the similarity of circuits across languages: a case study on the subject-verb agreement task](#). In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 10115–10125, Miami, Florida, USA. Association for Computational Linguistics.
- Matthew Finlayson, Aaron Mueller, Sebastian Gehrmann, Stuart Shieber, Tal Linzen, and Yonatan Belinkov. 2021. [Causal analysis of syntactic agreement mechanisms in neural language models](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1828–1843, Online. Association for Computational Linguistics.
- Dirk Goldhahn, Thomas Eckart, and Uwe Quasthoff. 2012. [Building large monolingual dictionaries at the leipzig corpora collection: From 100 to 200 languages](#). In *Proceedings of the Eighth International Conference on Language Resources and Evaluation (LREC 2012)*, Istanbul, Turkey. European Language Resources Association (ELRA).
- Hila Gonen, Yova Kementchedjhiava, and Yoav Goldberg. 2019. [How does grammatical gender affect noun representations in gender-marking languages?](#) In *Proceedings of the 2019 Workshop on Widening NLP*, pages 64–67, Florence, Italy. Association for Computational Linguistics.
- Phillip I. Good. 2005. *Permutation, Parametric, and Bootstrap Tests of Hypotheses*, 3 edition. Springer, New York.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. [The llama 3 herd of models](#). *arXiv preprint arXiv:2407.21783*.
- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). 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 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Matthew Honnibal, Ines Montani, Sofie Van Landeghem, and Adriane Boyd. 2020. [spaCy: Industrial-strength Natural Language Processing in Python](#).
- Yi Jing, Zijun Yao, Hongzhu Guo, Lingxu Ran, Xiaozhi Wang, Lei Hou, and Juanzi Li. 2025. [LinguaLens: Towards interpreting linguistic mechanisms of large language models via sparse auto-encoder](#). In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 28220–28239, Suzhou, China. Association for Computational Linguistics.
- Leipzig Corpora Collection. 2024. [German news corpus 2024 \(300k subset\)](#). https://downloads.wortschatz-leipzig.de/corpora/deu_news_2024_300k.tar.gz. Leipzig Corpora Collection. Dataset. Accessed: 2025-12-18.
- Jack Lindsey, Wes Gurnee, Emmanuel Ameisen, Brian Chen, Adam Pearce, Nicholas L. Turner, Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam Jermy, Andy Jones, and 8 others. 2025. [On the biology of a large language model](#). *Transformer Circuits Thread*.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. [Assessing the ability of LSTMs to learn syntax-sensitive dependencies](#). *Transactions of the Association for Computational Linguistics*, 4:521–535.
- Rebecca Marvin and Tal Linzen. 2018. [Targeted syntactic evaluation of language models](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1192–1202, Brussels, Belgium. Association for Computational Linguistics.
- Jan Pfister and Andreas Hotho. 2024. [SuperGLEBer: German language understanding evaluation benchmark](#). In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 7904–7923, Mexico City, Mexico. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, and 1 others. 2019.

Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.

Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. [A primer in BERTology: What we know about how BERT works](#). *Transactions of the Association for Computational Linguistics*, 8:842–866.

Wolfgang Seeker and Jonas Kuhn. 2013. [Morphological and syntactic case in statistical dependency parsing](#). *Computational Linguistics*, 39(1):23–55.

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.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). *Advances in neural information processing systems*, 30.

Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. [SuperGLUE: a stickier benchmark for general-purpose language understanding systems](#). Curran Associates Inc., Red Hook, NY, USA.

Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.

Alex Warstadt, Alicia Parrish, Haokun Liu, Anhad Mohananey, Wei Peng, Sheng-Fu Wang, and Samuel R. Bowman. 2020. [BLiMP: The benchmark of linguistic minimal pairs for English](#). *Transactions of the Association for Computational Linguistics*, 8:377–392.

Leonie Weissweiler, Valentin Hofmann, Anjali Kantharuban, Anna Cai, Ritam Dutt, Amey Hengle, Anubha Kabra, Atharva Kulkarni, Abhishek Vijayakumar, Haofei Yu, Hinrich Schuetze, Kemal Oflazer, and David Mortensen. 2023. [Counting the bugs in ChatGPT’s wugs: A multilingual investigation into the morphological capabilities of a large language model](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 6508–6524, Singapore. Association for Computational Linguistics.

Wikimedia Foundation. 2022. [Wikimedia wikipedia dataset](#). <https://huggingface.co/datasets/openskyml/wikipedia>. Version: “20220301.de”.

Julia Wunderle, Anton Ehrmanntraut, Jan Pfister, Fotis Jannidis, and Andreas Hotho. 2025. [New encoders](#)

for german trained from scratch: Comparing modernbert with converted llm2vec models. *Preprint*, arXiv:2505.13136.

A Data

Detailed data generation details for the article data and D_{NEUTRAL} , introduced in Section 4.

A.1 Article Data

This section provides details on generation for the GRADIEND training datasets, like $D_{\text{NOM}}^{\text{MASC}}$. Table 6 provides an overview of the generated datasets including sizes. Note that in this study, sometimes only a subset of the subsets is used for specific experiments, e.g., to balance datasets by min sampling. We release the dataset on Hugging Face under [aieng-lab/de-gender-case-articles](#).

A.1.1 Source Corpus and Sentence Segmentation

We use the German Wikipedia dump (snapshot 20220301.de) as the underlying text corpus (Wikimedia Foundation, 2022).

Due to the size of the German Wikipedia dump and the fact that the GRADIEND training does not require a very large number of data points, we do not process the corpus exhaustively. Instead, articles are subsampled using a fixed stride. Specifically, we extract short contiguous blocks of articles and skip large intervals between blocks. This yields a lightweight subset with broad topical coverage while avoiding locality effects introduced by processing consecutive articles only. The resulting subset is used solely as a source of naturally occurring sentences for morphosyntactic filtering and is not intended to represent a statistically uniform sample of Wikipedia.

A.1.2 Morphosyntactic Annotation

Each sentence is processed with spaCy to obtain token-level part-of-speech tags and morphological features. We rely on spaCy’s morphological annotations to identify the grammatical *case*, *gender*, and *number* of determiner tokens. Only tokens tagged as determiners (POS=DET) are considered as candidates for definite singular articles.

A.1.3 Sentence Filtering

For each gender-case combination $z = (g, c) \in \mathcal{G} \times \mathcal{C}$, we construct a separate dataset by retaining only sentences that satisfy all of the following constraints:

Dataset	ID	Size			
		Total	Train	Val	Test
$D_{\text{NOM}}^{\text{MASC}}$	masc_nom	34,350	27,829	3,084	3,437
$D_{\text{ACC}}^{\text{MASC}}$	masc_acc	30,538	24,781	2,705	3,052
$D_{\text{DAT}}^{\text{MASC}}$	masc_dat	23,437	18,918	2,176	2,343
$D_{\text{GEN}}^{\text{MASC}}$	masc_gen	34,087	27,417	3,254	3,416
$D_{\text{NOM}}^{\text{FEM}}$	fem_nom	61,328	49,399	5,796	6,133
$D_{\text{ACC}}^{\text{FEM}}$	fem_acc	34,801	28,155	3,166	3,480
$D_{\text{DAT}}^{\text{FEM}}$	fem_dat	46,601	37,458	4,482	4,661
$D_{\text{GEN}}^{\text{FEM}}$	fem_gen	38,811	31,219	3,711	3,881
$D_{\text{NOM}}^{\text{NEUT}}$	neut_nom	33,350	26,680	3,335	3,335
$D_{\text{ACC}}^{\text{NEUT}}$	neut_acc	19,012	15,209	1,901	1,902
$D_{\text{DAT}}^{\text{NEUT}}$	neut_dat	16,075	13,020	1,447	1,608
$D_{\text{GEN}}^{\text{NEUT}}$	neut_gen	25,351	20,436	2,387	2,528

Table 6: Dataset overview with full dataset sizes.

1. **Article presence.** The sentence contains at least one occurrence of the surface form corresponding to the target definite article.
2. **Morphological agreement.** All occurrences of the target article in the sentence are annotated with GENDER = g , CASE = c , and NUMBER = SING. Sentences containing plural uses of the article are excluded.
3. **Limited ambiguity.** Sentences containing more than four occurrences of the target article are discarded to reduce structural ambiguity.
4. **Length constraints.** Only sentences with a character length between 50 and 500 are retained.
5. **Named entity control.** Sentences containing more than three named entities are excluded to reduce confounds introduced by entity-heavy contexts.
6. **Duplicate removal.** Duplicate sentences are removed.

A.1.4 Data Quality

Error types. To quantify noise introduced by spaCy-based labeling, we manually annotated a small sample of instances from each gender-case dataset and analyzed the resulting errors. We distinguish two error categories:

- **In-paradigm cell errors (Cell):** spaCy assigns a token to the wrong gender-case cell *within* the definite singular article paradigm (e.g., $D_{\text{NOM}}^{\text{FEM}}$ mislabeled as $D_{\text{ACC}}^{\text{FEM}}$). These errors can be problematic because they effectively move examples

	Nom.	Acc.	Dat.	Gen.	All
Male	0.80	1.00	0.80	0.70	0.83
Neutral	0.90	0.85	0.95	0.70	0.85
Female	0.80	0.70	0.85	0.90	0.81
All	0.83	0.85	0.87	0.77	0.83

Table 7: Accuracy of spaCy-labeled gender-case cell data.

	$D_{\text{NOM}}^{\text{MASC}}$	$D_{\text{DAT}}^{\text{FEM}}$	$D_{\text{GEN}}^{\text{FEM}}$	Other
$D_{\text{NOM}}^{\text{MASC}}$	20	0	1	4
$D_{\text{DAT}}^{\text{FEM}}$	1	20	1	1
$D_{\text{GEN}}^{\text{NEUT}}$	0	1	23	1

Table 8: *der*: confusion matrix (spaCy rows; manual columns).

	$D_{\text{NOM}}^{\text{FEM}}$	$D_{\text{ACC}}^{\text{FEM}}$	Other
$D_{\text{NOM}}^{\text{FEM}}$	21	0	4
$D_{\text{ACC}}^{\text{FEM}}$	4	19	2

Table 9: *die*: confusion matrix (spaCy rows; manual columns).

between our datasets and may introduce misleading transition cues.

- **Out-of-paradigm errors (Other):** spaCy assigns a token to a category *outside* singular definite articles, e.g., relative pronouns, plural determiners, or other non-target uses. These instances primarily add *random* noise, since they do not systematically correspond to any other gender-case dataset used in our analyses.

Evaluation protocol. We manually labeled 20 randomly sampled text instances from each of the 12 gender-case datasets (Table 1). Labels correspond to the intended gender-case cell of the masked singular definite article (or *Other* if the token is not a singular definite article in context).

Overall accuracy. Table 7 reports accuracies aggregated by gender and case. spaCy achieves an overall accuracy of 0.83 across all datasets, with the lowest performance in the genitive cells (0.7). Accuracy is measured based on texts, i.e., texts containing multiple masks are only classified as True Positive if all its masks are correct.

Confusion matrices. To characterize *Cell* vs. *Other* errors in more detail, Tables 8–12 report confusion matrices for each surface form (excluding the non-syncretic *den* because of perfect accu-

	$D_{\text{NOM}}^{\text{NEUT}}$	$D_{\text{ACC}}^{\text{NEUT}}$	Other
$D_{\text{NOM}}^{\text{NEUT}}$	19	1	1
$D_{\text{NOM}}^{\text{NEUT}}$	3	17	0

Table 10: *das*: confusion matrix (spaCy rows; manual columns).

	$D_{\text{DAT}}^{\text{MASC}}$	$D_{\text{DAT}}^{\text{NEUT}}$	Other
$D_{\text{DAT}}^{\text{MASC}}$	18	3	1
$D_{\text{DAT}}^{\text{NEUT}}$	1	19	0

Table 11: *dem*: confusion matrix (spaCy rows; manual columns).

	$D_{\text{GEN}}^{\text{MASC}}$	$D_{\text{GEN}}^{\text{NEUT}}$	Other
$D_{\text{GEN}}^{\text{MASC}}$	15	7	1
$D_{\text{GEN}}^{\text{NEUT}}$	3	14	3

Table 12: *des*: confusion matrix (spaCy rows; manual columns).

racy). Rows correspond to spaCy assignments and columns to manual labels. Counts are *raw mask counts* (i.e., each mask is counted individually).

Discussion and implications for our analyses.

Across forms, *Other* errors are present, but they primarily inject random noise and are unlikely to induce systematic trends in gradient-based analyses. More critical are *Cell* errors within the article paradigm, which are most pronounced for genitive forms (notably *des*). Nevertheless, we observe that analyses relying on clean dataset combinations, in particular those involving *der* datasets and $D_{\text{NOM}}^{\text{FEM}}$ (*die*), match the conclusions obtained from noisier dataset combinations (see e.g., Figures 6, 8, and 36 with *clean* results). This consistency suggests that the overall trends reported in the paper are not driven solely by spaCy labeling noise.

A.2 Gender-Case Neutral Dataset (D_{NEUTRAL})

The D_{NEUTRAL} dataset is constructed from the Wortschatz Leipzig German news corpus¹ (Goldhahn et al., 2012; Leipzig Corpora Collection, 2024) to provide sentence contexts without grammatical gender or case cues. We apply a series of linguistic filters using spaCy to remove sentences that could implicitly encode such information.

Specifically, we exclude sentences that con-

¹https://downloads.wortschatz-leipzig.de/corpora/deu_news_2024_300K.tar.gz

tain determiners or definite and indefinite articles (including *der/die/das*, *ein/kein* and their inflected forms), as well as sentences containing third-person pronouns. To further reduce implicit gender signals, we remove sentences dominated by named entities, as proper names can carry gender information. We additionally filter out very short sentences and sentences containing the token *das* to avoid homonym-induced ambiguity.

The resulting dataset with 9,570 entries consists of well-formed sentences that are largely free of explicit and implicit morphosyntactic gender–case cues and is used as a grammar-neutral reference throughout our experiments.

We release the dataset on Hugging Face under [aieng-lab/wortschatz-leipzig-de-grammar-neutral](https://huggingface.co/aieng-lab/wortschatz-leipzig-de-grammar-neutral).

B Decoder-Only Models for MLM

Due to the fact that the gender of the definite article is usually only determined by the noun which naturally occurs in the right context, only using the left context is not an option for the considered problem, as done by (Drechsel and Herbold, 2026). Instead, we convert the decoder-only model into a model that can also predict a token given a right context, similar to the MLM task. We use the following general approach. We add a [MASK] token to the decoder’s tokenizer, and use the next $N > 0$ final hidden states of the decoder after the [MASK] token as mean-pooled input for a simple classifier network. The classifier has six classes, one for each German definite articles. This custom head makes it possible to use the decoder-only model for a bidirectional prediction task similar to a MLM task, at least to predict one of the six articles, and, importantly, create meaningful gradients through the entire model.

For the training of the classifier, we froze the core model parameters. This avoids changing the model, which could invalidate the GRADIEND models, as they may instead of analyzing the original model learn where the fine-tuning for this MLM-head updated the model. To choose an appropriate N , i.e., the hidden states after the [MASK] to consider as pooled classifier input, there is a natural trade-off. A low number might not include the encodings of the noun tokens (e.g., due to an adjective token(s) between the article and the noun). A too large number makes the relevant information from the noun less relevant, as it contributes less to

Model	# Parameters	Checkpoint	Reference
GermanBERT	109.1M	google-bert/bert-base-german-cased	Devlin et al. (2018)
GermanGPT-2	124.4M	dbmdz/german-gpt2	Radford et al. (2019)
EuroBERT	211.8M	EuroBERT/EuroBERT-210m	Boizard et al. (2025)
GBERT	335.7M	deepset/gbert-large	Chan et al. (2020)
ModernGBERT	1.06B	LSX-UniWue/ModernGBERT_1B	Wunderle et al. (2025)
LLaMA	3.21B	meta-llama/Llama-3.2-3B	Grattafiori et al. (2024)

Table 13: Hugging Face model checkpoints used in this study.

Hyperparameter	Value
Optimizer	Adam
Learning Rate	1×10^{-5} (encoder-only); 1×10^{-4} (decoder-only)
Weight Decay	1×10^{-2}
Batch Size Gradient Computation	4 (LLaMA), 32 (others)
Batch Size GRADIEND	1
Training Criterion	MSE
Training Steps	5,000
Evaluation Steps	1,000
Evaluation Max Size	500
Evaluation Criterion	Cor on validation split

Table 14: GRADIEND training hyperparameters.




Figure 10: Decoder-only article classifier performance across different pooling lengths.

the average pooling. We use $N = 3$ for GPT2 and $N = 5$ for LLaMA, as shown in Figure 10.

Overall, the classification performance is not too great considering the low number of classes (six), but sufficient to train the GRADIEND models.

C Training

Model details and training hyperparameters are reported in Tables 13 and 14, respectively. Following Drechsel and Herbold (2026), we train each GRADIEND variant with three random seeds and select the best run by the validation-set correlation as used in Table 3 (see Appendix D for details). For efficiency, we estimate this correlation on a 100-example subset per gender-case dataset from the validation split. We train LLaMA and ModernGBERT in `torch.bfloat16` and all other models in `torch.float32`.

We use an oversampled single-label batch sampler that groups examples by gender-case dataset

and constructs batches containing only one gender-case dataset label at a time. To ensure equal exposure across datasets, batches are oversampled to match the maximum number of batches per label and then interleaved in a round-robin fashion.

Experiments are run in Python 3.9.19. LLaMA is trained on three NVIDIA A100 GPUs (80 GB each), while all other models use a single A100. Per-seed training time ranges from ~ 1 hour (smaller models) to ~ 3 hours (LLaMA) for a single variant.

D Encoded Values

Figures 11–28 show encoded-value distributions for the remaining GRADIEND variants.

Correlations in Table 3 are computed from the same gradient types used during training. We assign labels $+1$ and -1 to the two directed transitions $z_1 \rightarrow z_2$ and $z_2 \rightarrow z_1$, respectively, and label the ten identity tasks $\tilde{z} \rightarrow \tilde{z}$ for $\tilde{z} \notin \{z_1, z_2\}$ as 0. Normalizing the sign of h during training makes correlations non-negative by construction (apart from extreme distribution shifts between the validation set used for normalization and the test set).

Since test splits differ in size across gender-case datasets, we compute correlations after randomly downsampling each dataset to the smallest test-set size. Likewise, for violins that combine multiple data sources in Figures 3 and 11–28, we downsample each source to the smallest subset contributing to that violin.

For D_{NEUTRAL} on decoder-only models, we use CLM gradients, since the auxiliary MLM head is restricted to predicting article tokens.

E Probability Analysis

Table 15 and Figures 29–34 show results for other models and variants reported in the main part of the paper. Due to computational constraints, we limit the German Language Understanding Evaluation Benchmark (SuperGLEBer) evaluation to models

Model	Art. Trans.	α	$D_{\text{NOM}}^{\text{MASC}}(\text{der})$			$D_{\text{GEN}}^{\text{FEM}}(\text{der})$			$D_{\text{DAT}}^{\text{FEM}}(\text{der})$			D_{NEUTRAL}			SuperGLEBer
			$\Delta\mathbb{P}$	d	Sig.	$\Delta\mathbb{P}$	d	Sig.	$\Delta\mathbb{P}$	d	Sig.	$\Delta\mathbb{P}$	d	Sig.	
GBERT	–	0.0	–	–	–	–	–	–	–	–	–	–	–	–	76.7 \pm 0.4
+ $G_{\text{NOM}}^{\text{FEM,MASC}}$	<i>der</i> \rightarrow <i>die</i>	0.05	0.13	0.27	***	0.01	0.11	***	0.02	0.14	***	-0.02	-0.00	n.s.	76.6 \pm 0.4
+ $G_{\text{ACC,DAT}}^{\text{FEM}}$	<i>der</i> \rightarrow <i>die</i>	1.0	0.70	0.36	***	0.45	0.34	***	4.20	0.71	***	0.12	0.03	*	76.6 \pm 0.4
+ $G_{\text{ACC,GEN}}^{\text{FEM}}$	<i>der</i> \rightarrow <i>die</i>	0.05	0.01	0.25	***	0.01	0.17	***	0.01	0.13	***	-0.06	-0.01	n.s.	76.6 \pm 0.4
+ $G_{\text{NOM,DAT}}^{\text{MASC}}$	<i>der</i> \rightarrow <i>dem</i>	0.01	0.00	0.12	***	0.00	0.08	***	0.00	0.04	*	-0.00	-0.00	n.s.	76.8 \pm 0.4
+ $G_{\text{FEM,NEUT}}^{\text{DAT}}$	<i>der</i> \rightarrow <i>dem</i>	0.01	0.00	0.08	***	0.00	0.04	***	0.00	0.16	***	0.02	0.01	n.s.	76.7 \pm 0.4
+ $G_{\text{NOM,GEN}}^{\text{MASC}}$	<i>der</i> \rightarrow <i>des</i>	0.01	0.00	0.08	***	0.02	0.22	***	0.00	0.07	***	-0.00	-0.01	n.s.	76.7 \pm 0.4
+ $G_{\text{FEM,NEUT}}^{\text{GEN}}$	<i>der</i> \rightarrow <i>des</i>	0.2	0.03	0.07	***	0.44	0.32	***	0.01	0.08	***	0.02	0.03	***	76.6 \pm 0.4
ModernGBERT	–	0.0	–	–	–	–	–	–	–	–	–	–	–	–	–
+ $G_{\text{NOM}}^{\text{FEM,MASC}}$	<i>der</i> \rightarrow <i>die</i>	0.2	10.67	0.69	***	3.09	0.38	***	7.28	0.52	***	1.26	0.14	***	–
+ $G_{\text{ACC,DAT}}^{\text{FEM}}$	<i>der</i> \rightarrow <i>die</i>	0.2	8.09	0.59	***	4.83	0.50	***	13.73	0.76	***	1.26	0.14	***	–
+ $G_{\text{ACC,GEN}}^{\text{FEM}}$	<i>der</i> \rightarrow <i>die</i>	0.001	0.01	0.19	***	0.00	0.11	***	0.01	0.16	***	0.02	0.00	n.s.	–
+ $G_{\text{NOM,DAT}}^{\text{MASC}}$	<i>der</i> \rightarrow <i>dem</i>	0.2	6.85	0.52	***	0.65	0.15	***	7.22	0.51	***	0.59	0.15	***	–
+ $G_{\text{FEM,NEUT}}^{\text{DAT}}$	<i>der</i> \rightarrow <i>dem</i>	0.2	7.15	0.47	***	2.17	0.23	***	27.06	1.02	***	0.83	0.17	***	–
+ $G_{\text{NOM,GEN}}^{\text{MASC}}$	<i>der</i> \rightarrow <i>des</i>	0.1	0.57	0.16	***	2.62	0.33	***	0.11	0.09	***	0.13	0.06	***	–
+ $G_{\text{FEM,NEUT}}^{\text{GEN}}$	<i>der</i> \rightarrow <i>des</i>	0.2	3.21	0.34	***	26.38	1.17	***	1.65	0.26	***	0.81	0.14	***	–
EuroBERT	–	0.0	–	–	–	–	–	–	–	–	–	–	–	–	68.3 \pm 0.4
+ $G_{\text{NOM}}^{\text{FEM,MASC}}$	<i>der</i> \rightarrow <i>die</i>	0.01	7.36	0.60	***	0.25	0.11	***	0.48	0.21	***	0.01	0.01	n.s.	68.4 \pm 0.4
+ $G_{\text{ACC,DAT}}^{\text{FEM}}$	<i>der</i> \rightarrow <i>die</i>	0.01	0.17	0.34	***	0.08	0.20	***	0.45	0.30	***	-0.00	-0.00	n.s.	67.0 \pm 0.4
+ $G_{\text{ACC,GEN}}^{\text{FEM}}$	<i>der</i> \rightarrow <i>die</i>	0.01	0.12	0.33	***	0.12	0.23	***	0.21	0.29	***	-0.00	-0.00	n.s.	67.5 \pm 0.4
+ $G_{\text{NOM,DAT}}^{\text{MASC}}$	<i>der</i> \rightarrow <i>dem</i>	0.01	0.16	0.32	***	0.00	0.09	***	0.03	0.18	***	-0.00	-0.00	n.s.	68.4 \pm 0.4
+ $G_{\text{FEM,NEUT}}^{\text{DAT}}$	<i>der</i> \rightarrow <i>dem</i>	0.01	0.18	0.17	***	0.06	0.05	***	1.28	0.28	***	0.00	0.04	***	67.4 \pm 0.4
+ $G_{\text{NOM,GEN}}^{\text{MASC}}$	<i>der</i> \rightarrow <i>des</i>	0.01	0.60	0.29	***	0.44	0.30	***	0.03	0.09	***	0.00	0.08	***	67.0 \pm 0.4
+ $G_{\text{FEM,NEUT}}^{\text{GEN}}$	<i>der</i> \rightarrow <i>des</i>	0.01	0.25	0.18	***	2.01	0.36	***	0.07	0.07	***	0.00	0.06	***	68.1 \pm 0.4
GermanGPT-2	–	0.0	–	–	–	–	–	–	–	–	–	–	–	–	45.4 \pm 0.4
+ $G_{\text{NOM}}^{\text{FEM,MASC}}$	<i>der</i> \rightarrow <i>die</i>	0.5	0.31	0.08	***	0.15	0.05	**	0.25	0.08	***	-0.00	-0.03	***	45.3 \pm 0.4
+ $G_{\text{ACC,DAT}}^{\text{FEM}}$	<i>der</i> \rightarrow <i>die</i>	0.5	0.20	0.04	*	0.25	0.09	***	0.24	0.07	***	0.00	0.06	***	45.3 \pm 0.4
+ $G_{\text{ACC,GEN}}^{\text{FEM}}$	<i>der</i> \rightarrow <i>die</i>	0.5	1.08	0.34	***	1.27	0.46	***	1.31	0.49	***	0.00	0.01	n.s.	45.3 \pm 0.4
+ $G_{\text{NOM,DAT}}^{\text{MASC}}$	<i>der</i> \rightarrow <i>dem</i>	0.1	0.48	0.48	***	0.10	0.20	***	0.06	0.04	**	0.00	0.09	***	45.3 \pm 0.4
+ $G_{\text{FEM,NEUT}}^{\text{DAT}}$	<i>der</i> \rightarrow <i>dem</i>	1.0	0.21	0.59	***	0.03	0.14	***	0.07	0.11	***	0.00	0.06	***	45.3 \pm 0.4
+ $G_{\text{NOM,GEN}}^{\text{MASC}}$	<i>der</i> \rightarrow <i>des</i>	0.1	0.47	0.27	***	0.49	0.11	***	0.53	0.19	***	0.00	0.02	*	45.3 \pm 0.4
+ $G_{\text{FEM,NEUT}}^{\text{GEN}}$	<i>der</i> \rightarrow <i>des</i>	0.5	0.43	0.23	***	1.39	0.27	***	0.61	0.20	***	0.00	0.02	**	45.3 \pm 0.4
LLaMA	–	0.0	–	–	–	–	–	–	–	–	–	–	–	–	–
+ $G_{\text{NOM}}^{\text{FEM,MASC}}$	<i>der</i> \rightarrow <i>die</i>	0.2	0.03	0.10	***	0.00	0.04	***	0.02	0.06	***	0.00	0.12	***	–
+ $G_{\text{ACC,DAT}}^{\text{FEM}}$	<i>der</i> \rightarrow <i>die</i>	0.2	0.03	0.10	***	0.00	0.04	***	0.03	0.07	***	0.00	0.11	***	–
+ $G_{\text{ACC,GEN}}^{\text{FEM}}$	<i>der</i> \rightarrow <i>die</i>	0.2	0.04	0.10	***	0.00	0.04	***	0.03	0.07	***	0.00	0.11	***	–
+ $G_{\text{NOM,DAT}}^{\text{MASC}}$	<i>der</i> \rightarrow <i>dem</i>	0.2	0.01	0.09	***	0.00	0.07	***	0.02	0.09	***	0.00	0.06	***	–
+ $G_{\text{FEM,NEUT}}^{\text{DAT}}$	<i>der</i> \rightarrow <i>dem</i>	0.5	0.03	0.09	***	0.00	0.07	***	0.09	0.08	***	0.01	0.11	***	–
+ $G_{\text{NOM,GEN}}^{\text{MASC}}$	<i>der</i> \rightarrow <i>des</i>	0.2	0.00	0.01	n.s.	0.03	0.08	***	0.00	0.03	*	0.00	0.05	***	–
+ $G_{\text{FEM,NEUT}}^{\text{GEN}}$	<i>der</i> \rightarrow <i>des</i>	0.2	0.01	0.25	***	0.09	0.15	***	0.02	0.12	***	0.00	0.07	***	–

Table 15: GRADIEND-modified models for non-GermanBERT models: mean change in target-article probability $\Delta\mathbb{P}$ scaled by 100, effect size (Cohen’s d), and significance as *** $p < .001$, ** $p < .01$, * $p < .05$ (n.s. otherwise). Bold marks datasets being part of GRADIEND gender-case cells. SuperGLEBer score is scaled by 100, and only reported for small models ($< 1\text{B}$).

with fewer than 1B parameters, and therefore exclude ModernGBERT and LLaMA.

E.1 Details on Probability Calculation

For an entry x from a fixed gender-case dataset (clear from context) with a single article mask, let $\mathbb{P}_m(\text{art}|x)$ denote the MLM/CLM probability that model m assigns to the article $\text{art} \in \mathcal{A}$ at the masked position for model m (treating tokens as equal up to casing and leading whitespace). We define the mean article probability $\mathbb{P}_m(\text{art})$ as the dataset average of $\mathbb{P}_m(\text{art}|x)$ over all single mask entries of the dataset. For two models m_1, m_2 , we

define the mean probability change as

$$\Delta\mathbb{P}_{m_1, m_2}(\text{art}) = \mathbb{P}_{m_1}(\text{art}) - \mathbb{P}_{m_2}(\text{art}).$$

For this study, m_1 is the α^* selected GRADIEND-modified model (clear from context) and m_2 the base model, so we write $\Delta\mathbb{P}(\text{art})$ or simply $\Delta\mathbb{P}$ when the article is clear from context.

E.2 Selection of the Intervention Strength α .

We evaluate a discrete set of step sizes $\alpha > 0$ that induce GRADIEND-modified models at α (indicated in Figure 5). Let s_0 denote the LMS of the unmodified base model on the grammar-neutral dataset D_{NEUTRAL} . For encoder-only models, this

score corresponds to masked-token accuracy, while for decoder-only models it corresponds to perplexity (lower is better).

We define a tolerance threshold $\tau = 0.99$ and select α^* according to the following procedure. Among all evaluated step sizes, we retain those as *candidates* whose score satisfies the constraint $s(\alpha) \geq \tau \cdot s_0$ for accuracy-based metrics or $s(\alpha) \leq s_0/\tau$ for perplexity-based metrics. This candidate range is shaded in Figure 5. Among candidates, we select the α , whose score is the largest: $\alpha^* = \operatorname{argmax} \mathbb{P}_\alpha(\text{target})$.

E.3 SuperGLEBer

We evaluate downstream language understanding using the German Language Understanding Evaluation Benchmark (SuperGLEBer; Pfister and Hotho 2024), a German NLP benchmark covering multiple classification and inference tasks. We follow the standard implementation provided by Pfister and Hotho (2024).

We note that for some model-task combinations, the Named Entity Recognition (NER) components yield *empty entity predictions*, resulting in zero precision and recall, a known degeneracy in sequence labeling. We attribute this to a dependency incompatibility in our evaluation pipeline, since it occurs consistently for both the base model and the corresponding GRADIEND-modified variants.

We extend their evaluation code to support bootstrap-based uncertainty estimation (Davison and Hinkley, 1997). Our procedure mirrors the approach used by Drechsel and Herbold (2026) for GLUE (Wang et al., 2018) and SuperGLUE (Wang et al., 2019), computing 95% confidence intervals via resampling.

F Top- k Analysis

Figures 35–37 show the Venn diagrams of non-GermanBERT-models similar to Figures 7–9.

Our overlap analysis depends on the choice of k . Choosing k too small makes overlaps sensitive to ranking noise and a few extreme weights, whereas choosing k too large gradually includes many weakly-informative weights, making overlap less diagnostic.

To motivate our choice, we ablate k and plot pairwise Top- k overlap within each article group as a function of k (Figures 38–40). Across model/group combinations, the curves vary in detail, but many exhibit a recurring three-range pattern:

- (i) **Transition-dominated range (small/intermediate k):** overlap is comparatively high and relatively stable, indicating that the Top- k sets are dominated by weights that are consistently ranked highly across variants within the same article group.
- (ii) **Reduction range (intermediate/large k):** overlap often decreases as increasing k begins to include additional weights whose ranks are less consistent across variants, reducing the intersection proportion.
- (iii) **Trivial convergence range (very large k):** overlap increases again as selectivity vanishes; in the limit, overlap approaches 100% when k becomes large relative to the parameter count.

While this trend is not uniform for all models and variants, $k=1000$ yields a stable and comparable operating point across article groups, focusing on a small but informative subset of weights. We choose $k=1000$ because it lies in the first, transition-dominated regime for most model-group combinations in Figures 38 and 39, before the dilution-driven decrease becomes prominent. Importantly, the intersection proportions of the article groups (Figures 38 and 39) are consistently above the control group proportions (Figure 40) for $k \leq 10,000$, indicating a generalization of our main claim based on Table 5 across a wide range of small k .




Figure 11: Encoded value distribution of $G_{NOM,DAT}^{FEM}$ for different input gradients.




Figure 13: Encoded value distribution of $G_{ACC,DAT}^{FEM}$ for different input gradients.




Figure 12: Encoded value distribution of $G_{NOM,GEN}^{FEM}$ for different input gradients.

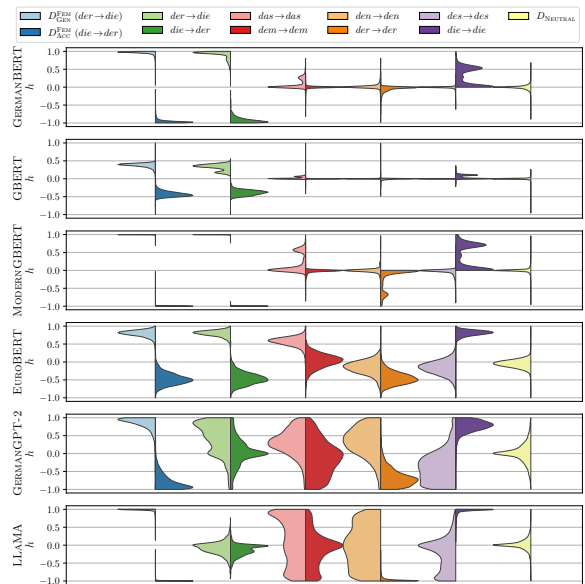


Figure 14: Encoded value distribution of $G_{ACC,GEN}^{FEM}$ for different input gradients.




Figure 15: Encoded value distribution of $G_{\text{DAT}}^{\text{FEM,MASC}}$ for different input gradients.




Figure 17: Encoded value distribution of $G_{\text{NOM,DAT}}^{\text{MASC}}$ for different input gradients.




Figure 16: Encoded value distribution of $G_{\text{DAT}}^{\text{FEM,NEUT}}$ for different input gradients.




Figure 18: Encoded value distribution of $G_{\text{GEN}}^{\text{FEM,MASC}}$ for different input gradients.




Figure 19: Encoded value distribution of $G_{\text{GEN}}^{\text{FEM,NEUT}}$ for different input gradients.

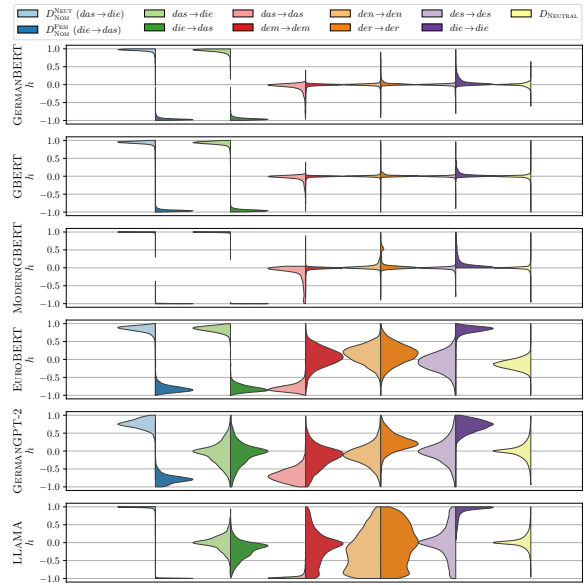


Figure 21: Encoded value distribution of $G_{\text{NOM}}^{\text{FEM,NEUT}}$ for different input gradients.




Figure 20: Encoded value distribution of $G_{\text{NOM,GEN}}^{\text{MASC}}$ for different input gradients.




Figure 22: Encoded value distribution of $G_{\text{ACC}}^{\text{FEM,NEUT}}$ for different input gradients.

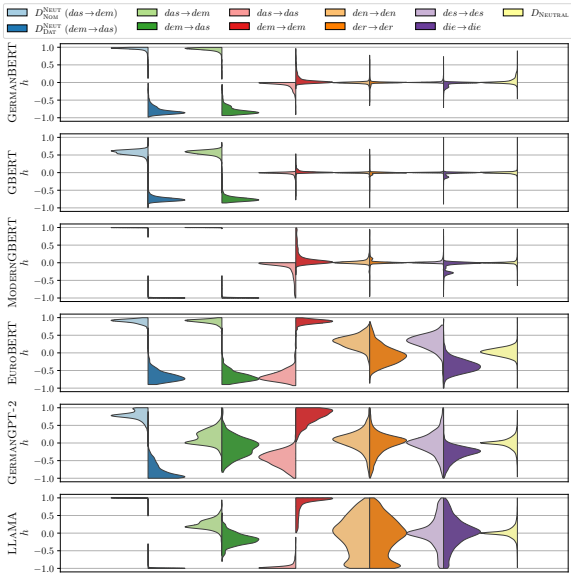


Figure 23: Encoded value distribution of $G_{\text{NOM, DAT}}^{\text{NEUT}}$ for different input gradients.




Figure 25: Encoded value distribution of $G_{\text{NOM, GEN}}^{\text{NEUT}}$ for different input gradients.

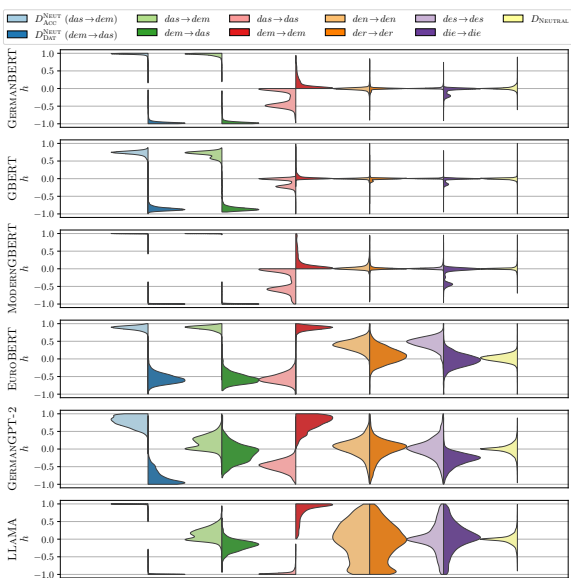


Figure 24: Encoded value distribution of $G_{\text{ACC, DAT}}^{\text{NEUT}}$ for different input gradients.

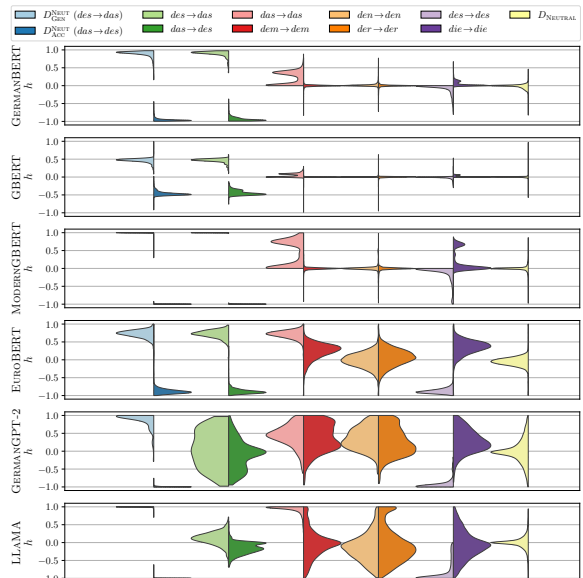


Figure 26: Encoded value distribution of $G_{\text{ACC, GEN}}^{\text{NEUT}}$ for different input gradients.




Figure 27: Encoded value distribution of $G_{\text{GEN,DAT}}^{\text{MASC}}$ for different input gradients.




Figure 28: Encoded value distribution of $G_{\text{GEN,DAT}}^{\text{NEUT}}$ for different input gradients.

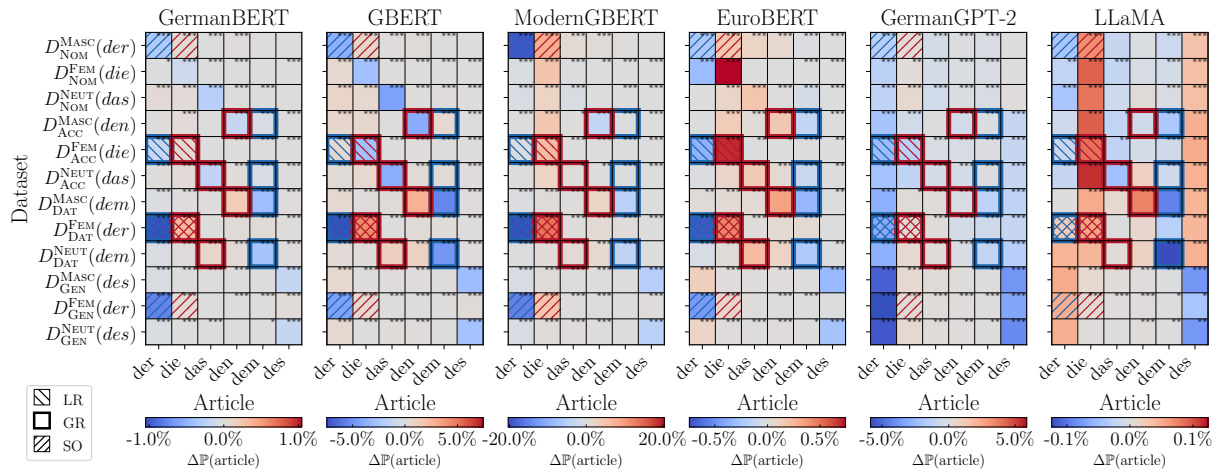


Figure 29: Mean probability change of articles between GRADIEND-modified and base model for $G_{\text{ACC,DAT}}^{\text{FEM}} \text{der} \rightarrow \text{die}$. Stars mark statistical significance after Benjamini-Hochberg FDR correction applied per model. Marked cells are expectations for LR, GR, and SO (Figure 4).

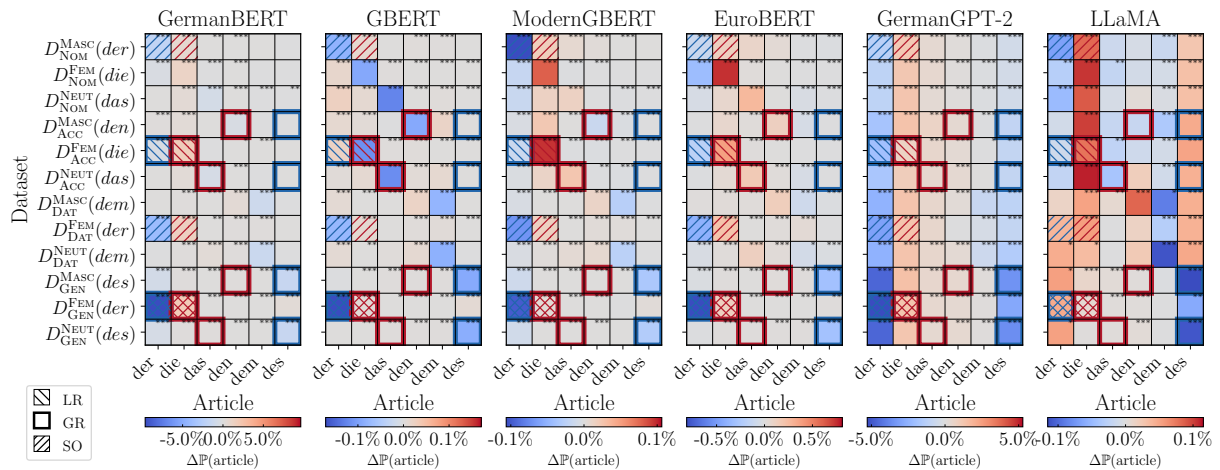


Figure 30: Mean probability change of articles between GRADIEND-modified and base model for $G_{\text{ACC,GEN}}^{\text{FEM}} \text{der} \rightarrow \text{die}$. Stars mark statistical significance after Benjamini-Hochberg FDR correction applied per model. Marked cells are expectations for LR, GR, and SO (Figure 4).

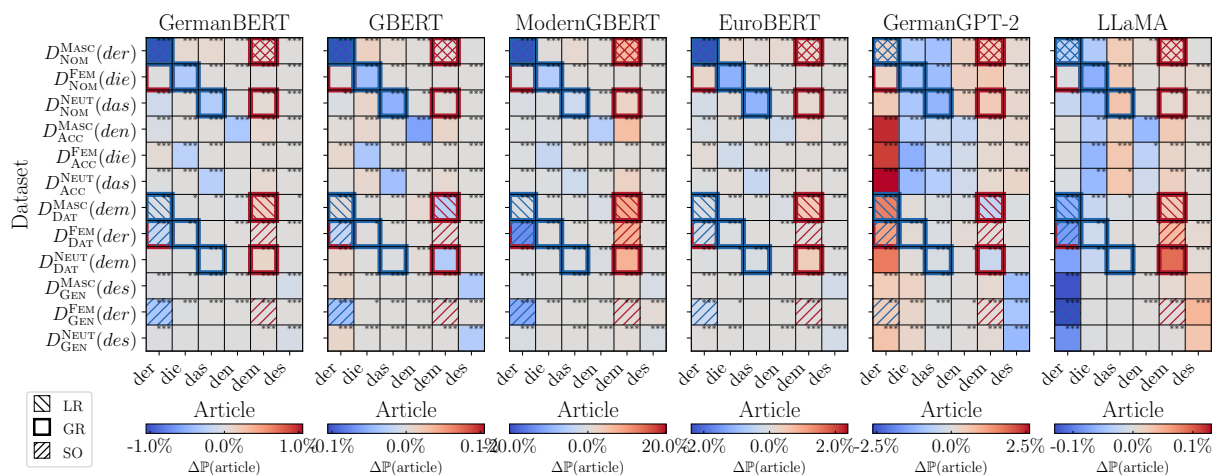


Figure 31: Mean probability change of articles between GRADIEND-modified and base model for $G_{\text{NOM,DAT}}^{\text{MASC}} \text{der} \rightarrow \text{dem}$. Stars mark statistical significance after Benjamini-Hochberg FDR correction applied per model. Marked cells are expectations for LR, GR, and SO (Figure 4).

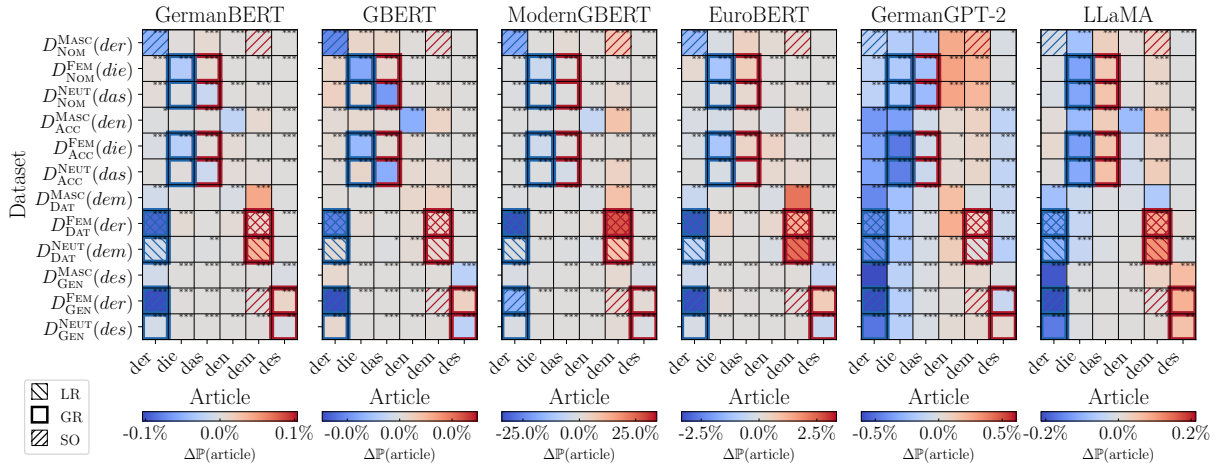


Figure 32: Mean probability change of articles between GRADIEND-modified and base model for $G_{\text{DAT}}^{\text{FEM, NEUT}}$ $der \rightarrow dem$. Stars mark statistical significance after Benjamini-Hochberg FDR correction applied per model. Marked cells are expectations for LR, GR, and SO (Figure 4).

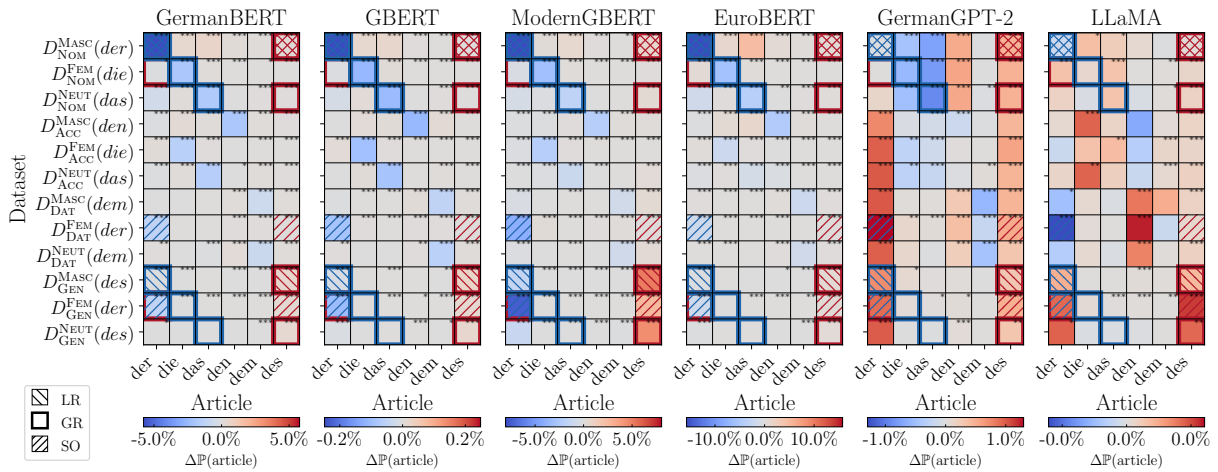


Figure 33: Mean probability change of articles between GRADIEND-modified and base model for $G_{\text{NOM, GEN}}^{\text{MASC}}$ $der \rightarrow des$. Stars mark statistical significance after Benjamini-Hochberg FDR correction applied per model. Marked cells are expectations for LR, GR, and SO (Figure 4).

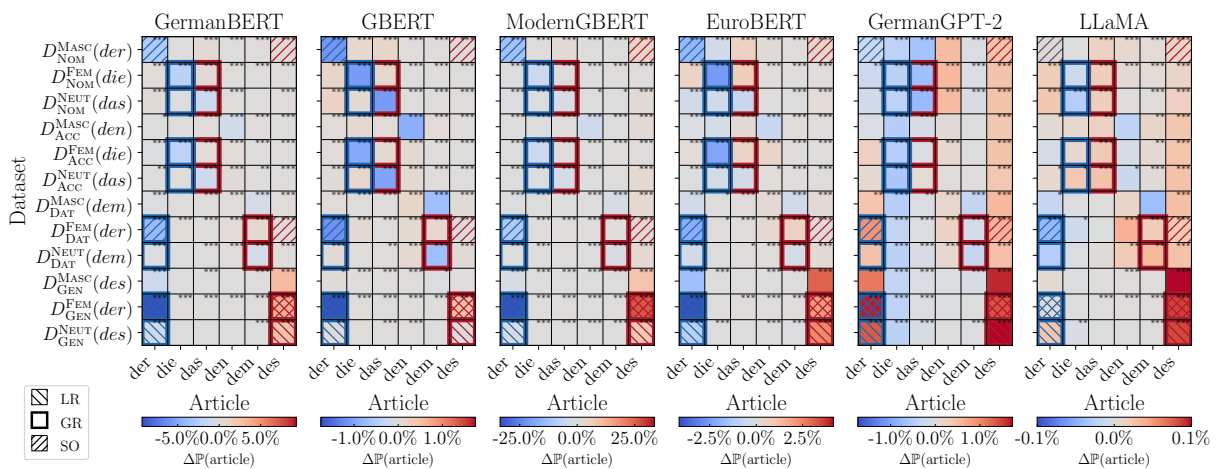


Figure 34: Mean probability change of articles between GRADIEND-modified and base model for $G_{\text{GEN}}^{\text{FEM, NEUT}}$ $der \rightarrow des$. Stars mark statistical significance after Benjamini-Hochberg FDR correction applied per model. Marked cells are expectations for LR, GR, and SO (Figure 4).

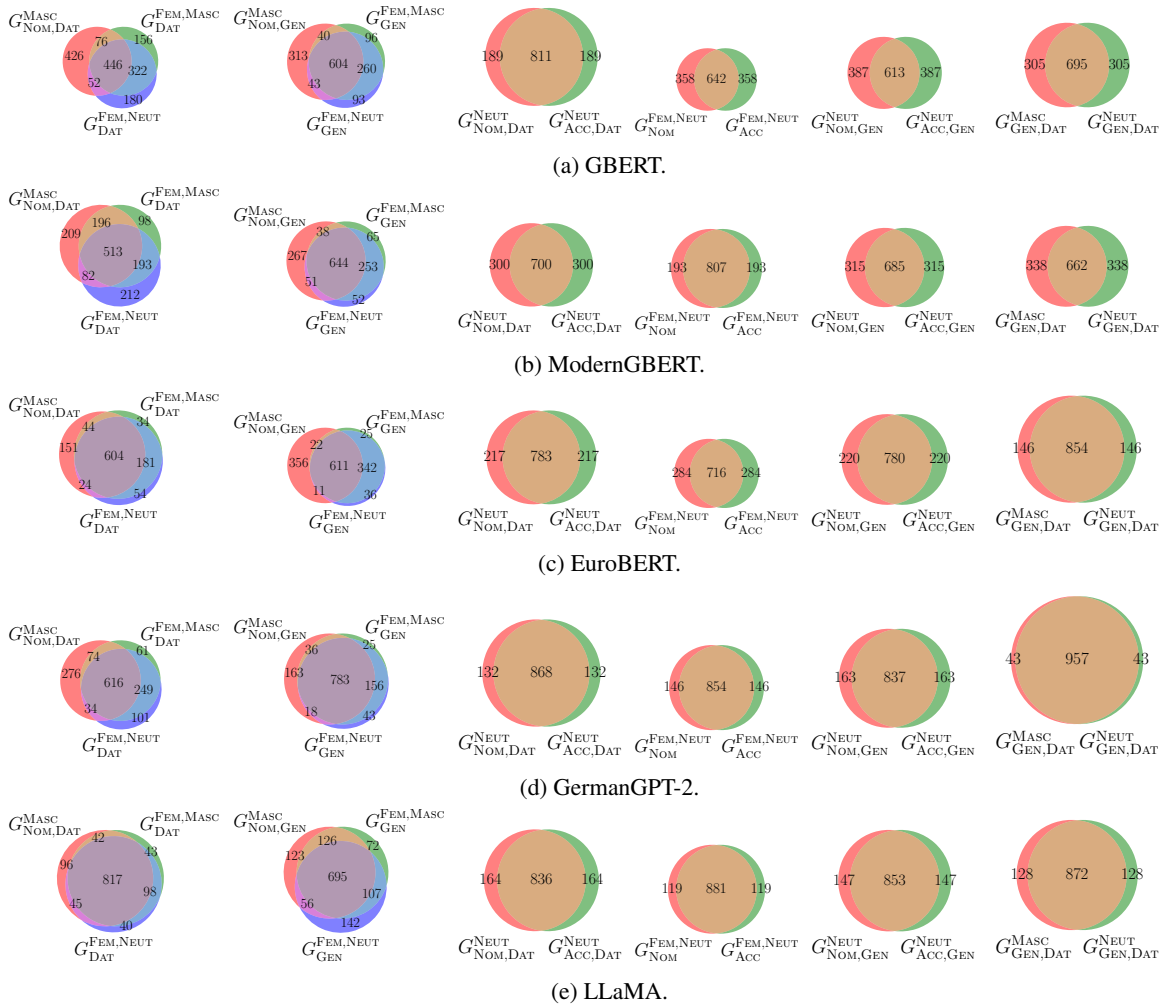


Figure 35: Top-1,000 weight overlaps across different GRADIENDs for non-GermanBERT-models.




Figure 36: Top-1,000 weight overlaps for non-GermanBERT models of article group *der* ↔ *die*.




Figure 37: Top-1,000 weight overlaps for non-GermanBERT models of the control group.

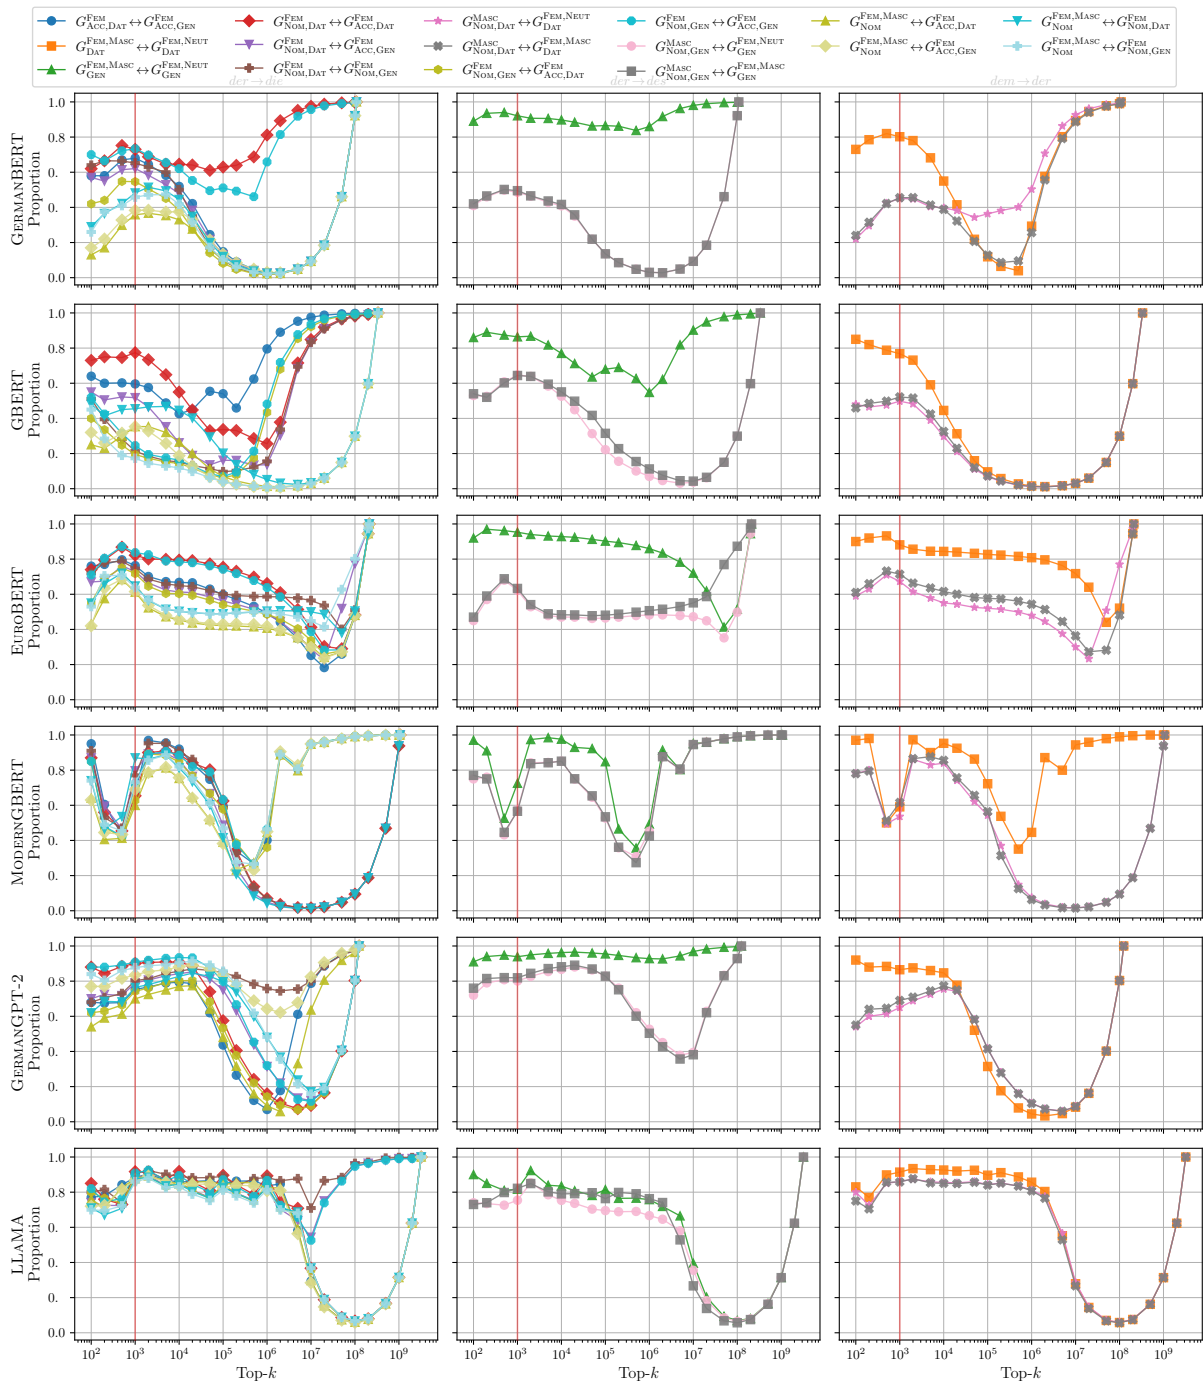


Figure 38: Pairwise Top- k weight-intersection proportions across values of k for the control group for the two-dimensional article groups.




Figure 39: Pairwise Top- k weight-intersection proportions across values of k for the one-dimensional article groups.




Figure 40: Pairwise Top- k weight-intersection proportions across values of k for the control group.