

Beyond Lexical and Syntactic: Better Text Simplicity Analysis Informed by Psycholinguistic Research

Benjamin Vendeville^{1,2} , Liana Ermakova^{1,3} , Pierre De Loor^{2,4} 

¹Université de Bretagne Occidentale, Brest, France

²Lab-STICC (UMR CNRS 6285), Brest, France, ³HCTI, Brest, France, ⁴ENIB, Brest, France

Correspondence: benjamin.vendeville@univ-brest.fr, liana.ermakova@univ-brest.fr, deloor@enib.fr

Abstract

Automatic text simplification measures correlate poorly with human judgments and are disconnected from psycholinguistic research on comprehension. We propose a multidimensional analysis framework grounded in Cognitive Load Theory and the Construction-Integration model, organizing 93 construct-validated linguistic features (selected from a starting pool of over 400) into six simplicity dimensions corresponding to stages of the comprehension process as described in the literature. Using this framework, we analyze four human-annotated simplicity datasets and six aligned simplification splits. We find that (1) our dimensions are informative enough to predict human simplicity ratings on most datasets; (2) current datasets overwhelmingly target lexical and syntactic dimensions while under-targeting personalization, the dimension psychological research identifies among the most impactful for comprehension; and (3) existing measures overall do not capture any specific dimension. These findings provide actionable guidance for building measures and datasets better aligned with how readers actually process text.

1 Introduction

Despite rising global literacy, many texts remain inaccessible to general audiences due to their complexity. Automatic Text Simplification (ATS) addresses this with systems trained on corpora of aligned complex-simple text pairs (Saggion, 2017), increasingly Large Language Models (LLMs) (Ermakova et al., 2026). These models are then evaluated on similar datasets, using various measures that compare the generated text to reference simplified versions such as SARI (the most widely used ATS measure) (Xu et al., 2016), BLEU (Papineni et al., 2002), or without references like FKGL (Kincaid et al., 1975). However, these measures have repeatedly shown low correlation with

human ratings of *simplicity* (Alfear et al., 2024; Alva-Manchego et al., 2021), but also with *meaning preservation* (Agrawal and Carpuat, 2024) or *hallucination* (Vendeville et al., 2025). They are also often disconnected from findings in psycholinguistic and cognitive research, which highlight that simplicity depends on multiple interacting dimensions such as cohesion, coherence or even personalization (Strohmaier et al., 2023). Moreover, both measures and datasets tend to focus on one or two dimensions of simplicity (typically lexical and syntactic simplicity) (Davari et al., 2024), producing single isolated scores. These offer limited insight into what is actually happening during simplification, making it difficult to reliably evaluate ATS.

We argue that psycholinguistic research can inform ATS evaluation in a more human-correlated direction. We therefore propose a multidimensional framework that operationalizes dimensions identified as salient for comprehension in the broader psycholinguistic literature. We draw on Cognitive Load Theory (CLT) (Sweller, 1988; Sweller et al., 1998) and the Construction-Integration (CI) model (Kintsch, 1998) for the theoretical grounding. The operational dimensions and their definitions are taken from the meta-analytic coding scheme of Strohmaier et al. (2023). From this research, we assign text features to dimensions based on whether their construct-validated definitions from the source tools match the paper’s dimensional definitions. We use this framework to identify which dimensions of simplicity current measures and datasets capture, and which they miss. Specifically, we address three research questions: **RQ1**: Do psycholinguistically-grounded simplicity dimensions predict human simplicity judgments, and does their relative importance vary across datasets? We train per-dimension regressors on 4 annotated datasets and compare their predictive power to established measures. **RQ2**: Which simplicity dimensions do different simplification

datasets actually transform during simplification and does the magnitude change between datasets? We measure per-dimension effect sizes ($|\Delta z|$) on 6 aligned datasets, validated by per-dimension binary classification. **RQ3:** Which simplicity dimensions do existing evaluation measures capture? We correlate measure scores with dimension-level predictions across datasets to reveal measure biases. All code and data are released as a Python library, SIMPSY (Appendix H).

2 Related Work

2.1 Limitations of Current ATS Measures

Evaluation of text simplification requires measures that can both predict human judgments, as well as be interpreted to guide system development. Current methods fall short on both. Indeed, most measures produce a single score that conflates distinct aspects of simplicity (often just lexical difficulty and syntactic complexity) into one number, making it impossible to diagnose *why* a simplification succeeds or fails. This is true of classical readability formulas (Kincaid et al., 1975; Flesch, 1948), which reduce comprehension difficulty to word and sentence length; or reference-based measures like SARI (Xu et al., 2016), BLEU (Papineni et al., 2002), and ROUGE (Lin, 2004), which measure textual similarity without distinguishing what kind of linguistic transformation was applied. Similarly, neural approaches like LENS (Maddela et al., 2023), BERTScore (Zhang et al., 2020), and SLE (Cripwell et al., 2023) can improve predictive power but remain opaque. Across all these classes, correlation with human ratings of simplicity is inconsistent and at best moderate (Alfear et al., 2024; Alva-Manchego et al., 2021; Cardon and Doğruöz, 2025).

Beyond simplicity, Agrawal and Carpuat (2024) and Davari et al. (2024) demonstrate that standard measures largely fail to capture meaning preservation, on reading-comprehension data and across several simplification corpora respectively. Vendeville et al. (2025) show that faithfulness measures (FactCC, QAGS, FEQA) transfer poorly to simplification data, and propose a detailed error taxonomy together with an annotated test collection. These failures point to a structural problem rather than a calibration one. Simplicity is not a single construct: psycholinguistic research shows that comprehension depends on the interaction of multiple dimensions such as lexical, syntactic, but also others such

as discourse, and prior knowledge (Kintsch, 1998; Ozuru et al., 2009). Despite that, no existing measure decomposes its score along these lines. Without such decomposition, measures cannot explain domain-specific failures or provide actionable feedback for improving simplification systems.

2.2 Toward Multidimensional Assessment

A few approaches have moved beyond single-score evaluation. TextEvaluator (Sheehan et al., 2014) combines 43 features across lexical, syntactic, and discourse categories, but uses mainly closed-source tools such as word dictionaries. BATS (Kreutz et al., 2024) implements 37 binary simplicity and complexity characteristics from the literature and analyzes their predictive importance across domains and audiences. They notably show that feature relevance varies substantially by context.

However, these approaches share critical limitations. They identify predictive features without organizing them into theoretically coherent dimensions. BATS observes a domain-dependent variation but does not explain *why* certain features matter more in certain contexts. Neither framework treats *discourse coherence* as a distinct dimension, although cohesion and coherence are long-established determinants of comprehension (McNamara and Kintsch, 1996; Ozuru et al., 2009). TextEvaluator does include a cohesion component but limits it to content-word overlap and explicit connectives (Sheehan et al., 2014); BATS includes some entity-tracking features that overlap with cohesion but does not organize them as a distinct dimension (Kreutz et al., 2024). Most importantly, they lack grounding in cognitive theory that could provide interpretability, by connecting observable text features to the processing stages and cognitive loads they affect during reading.

Recent meta-evaluation work by Cardon and Doğruöz (2025) showed this gap. They show weak and inconsistent correlations between readability measures, automatic metrics, and human judgments on English ATS datasets. They attribute this to the lack of a well-defined construct in the field. Our framework is one response to this gap: we derive simplicity dimensions from the Construction-Integration model (Kintsch, 1998) and Cognitive Load Theory (Sweller, 1988). Then, rather than aggregating features into a single, opaque score, we define multiple simplicity dimension. Each one is grounded in a specific comprehension mechanism. Appendix B reports a detailed comparison with Tex-

tEvaluator and BATS along on feature provenance, dimension construction, and theoretical grounding.

3 Methodology

We ground our analysis in psycholinguistic theory (Section 3.1) to derive six simplicity dimensions (Section 3.3), which we operationalize using 445 text features. Using this framework, we examine whether these dimensions predict simplicity ratings (**RQ1**); which dimensions are transformed during simplification and by how much (**RQ2**); and which dimensions existing evaluation measures capture (**RQ3**). We evaluate across eight corpora spanning different domains, granularities, and target audiences. Those include four *simplicity* datasets, which provide per-sample human simplicity ratings, and four *simplification* datasets, which provide aligned complex–simple pairs without ratings (six aligned splits in total, since OneStopEnglish contributes three level pairs). Specifically, we include Simplicity-DA (Alva-Manchego et al., 2021) (SDA), which rates system outputs on TurkCorpus sentences (Xu et al., 2016), and D-Wiki (Sun et al., 2021), built from English and Simple English Wikipedia, at the sentence and document level respectively. We also include the SimpleText dataset (Ermakova et al., 2024), which is built on manually written simplifications of scientific abstracts, and includes a sentence level (ST_{sent}) and a paragraph level (ST_{para}) split. All of those datasets are provided with human ratings of simplicity for each sample. For D-Wiki these per-sample ratings are the ones distributed with the dataset release, rather than the smaller human evaluation reported in the paper. In the case of ST, an unspecified number of writers explicitly provided evaluations of *lexical simplicity* and *syntax simplicity* (Davari et al., 2024), which we averaged to get a single simplicity score. We also include four simplification datasets that provide no per-sample simplicity ratings for the aligned pairs: ASSET (Alva-Manchego et al., 2020), a multi-reference corpus of English Wikipedia sentences rewritten with several simplification operations, Cochrane (Devaraj et al., 2021), a dataset of medical abstracts simplified by experts, Cochrane_{auto} (Bakker and Kamps, 2024), a re-aligned version of Cochrane, and finally OneStopEnglish (Vajjala and Lučić, 2018) (OSE), a dataset composed of educational texts written at three levels Elementary, Intermediate, and Advanced (respectively OSE_E, OSE_I and OSE_A). Ta-

Dataset	Samples	Domain	Align	Level	Raters
SDA	600	Wiki	C,S,R	s	15
ST _{sent, para}	1.6K, 271	Science	C,S,R	s,m	N/A
D-Wiki	601	Wiki	C,S,R	d	3
ASSET	4,500	Wiki	C,S,R	s	-
OneStopEnglish	2663	Edu.	E,I,A	m	-
Cochrane	3967	Medical	C,S	d	-
Cochrane _{auto}	849	Medical	C,S	d	-

Table 1: Dataset properties. Align. shows available texts: C=Complex, S=Simple, R=Reference. Level shows text granularity: s=sentence, m=multi-sentence (paragraph), d=document. Number of raters not provided for SimpleText data; ASSET, OSE and the Cochrane datasets carry no simplicity ratings.

ble 1 summarizes those datasets.

3.1 Cognitive Models

We use two frameworks: the Construction-Integration (CI) model (Kintsch, 1998), which describes the phases of text comprehension, and Cognitive Load Theory (CLT) (Sweller, 1988; Sweller et al., 1998), which models how limited cognitive resources are allocated, differently across readers (Ozuru et al., 2009). Figure 1 summarizes how these map to our simplicity dimensions.

In the CI model, reading comprehension is presented as a three-stage process: (1) *construction*, where readers identify concepts from the text, (2) *integration*, where readers link identified concepts into a textbase graph representing the text, and (3) *situation model construction* where the reader integrates the text representation with prior knowledge into a coherent mental model called a *situational model*.

The CLT models task difficulty as a function of working-memory effort. Working memory is the limited-capacity system used to process and maintain information during cognitive tasks. Tasks impose cognitive loads that consume this capacity. It distinguishes three forms of cognitive load: *intrinsic load*, determined by the inherent complexity of the content or task; *extraneous load*, which is the load imposed by how information is presented (e.g., linguistic form, discourse structure, layout); and *germane load*, in contrast, represents the desirable cognitive effort dedicated to constructing and refining mental representations, and is therefore essential for meaningful comprehension.




Figure 1: Unified Simplicity Framework. The overhead arrow indicates Reading Comprehension Steps through the three Construction-Integration stages. Cognitive Load Theory operates in parallel, characterizing load types at each stage. Linguistic and Content simplicity dimensions target features that modulate these loads; Personalization operates separately through motivation and engagement.

3.2 Text Comprehension Stages

In [Strohmaier et al. \(2023\)](#), a meta-analysis of 45 studies with 6,477 learners, the authors examine how linguistic modifications impact comprehension and learning from STEM texts. We group their categories of linguistic modification onto the three CI processes described above: reducing complexity supports the construction stage, increasing within-text links supports integration, and increasing links with prior knowledge supports situation model construction. The cognitive load attribution to each follows from CLT. Figure 1 summarizes the multidimensional nature of comprehension and illustrates how different cognitive stages and loads map to potential ATS interventions.

During the *construction* stage, readers decode words and linguistic units, constructing propositions from lexical and syntactic forms. This process is heavily influenced by text presentation and induces high *extraneous load*. Word frequency modulates lexical processing effort ([Berglund-Barraza et al., 2019](#)), and reducing word-level complexity (scientific vocabulary...) or syntactic complexity (sentence length, syntax simplification...) is intended to minimize this load. However, [Strohmaier et al. \(2023\)](#) shows that such surface-level complexity reduction shows negligible or inconsistent effects on comprehension and learning.

During the *integration* stage, readers link con-

cepts into a connected textbase representing the text. This process contributes moderately to *germane load*. Increasing cohesion, referential overlap, and connectives aim to reduce *extraneous load* by making semantic connections explicit ([McNamara and Kintsch, 1996](#)). Yet cohesion modifications alone show weak or null average effects ([Strohmaier et al., 2023](#)), and where they do help, the benefit is moderated by reader knowledge and comprehension skill ([O'Reilly and McNamara, 2007](#)) and by executive function ([Follmer and Sperling, 2018](#)), suggesting that just freeing working memory capacity provides minimal benefit. Importantly, prior knowledge plays an important part here. For expert readers, reducing extraneous load at this stage might bring no improvement, or even hinder comprehension ([Rey and Buchwald, 2011](#)), in what is called *expertise reversal effect*. This effect is not symmetrical: the meta-analysis finds that additional assistance helps novices more than it hurts experts ([Tetzlaff et al., 2025](#)).

During the *situation model construction*, readers integrate the text representation with prior knowledge to build a coherent mental model ([O'Reilly and McNamara, 2007](#); [Graesser et al., 2011](#)). This stage depends heavily on both *germane load* (productive effort) and *intrinsic load* (inherent content difficulty) ([Strohmaier et al., 2023](#)). Increasing clarity and elaboration (by explaining concepts,

specifying causal relations, personalization etc...) helps germane load. Notably, these modifications show significant positive effects. Finally, conceptual density (number of concepts mentioned per unit of text) and concept abstractness increase the intrinsic load and significantly influence comprehension difficulty.

We follow the categories and definitions from [Strohmaier et al. \(2023\)](#) for our framework. We keep dimensions regardless of if they were found important for simplicity in their context, RQ2 examines how dimensional importance varies across populations and text types. We distinguish two main categories: *Linguistic simplicity* and *Content simplicity*. Linguistic simplicity is composed of modifications that mainly reduce extraneous load across comprehension stages by transforming the actual words being used and their order. Content simplicity, in contrast, addresses mainly intrinsic load by managing the concepts being used and their interactions.

3.3 Simplicity Dimensions

We derive six simplicity dimensions (summarized in Table 2) anchored on the dimensional definitions in [Strohmaier et al. \(2023\)](#). We operationalize each dimension using subsets of 445 linguistic features extracted from publicly available tools and resources, listed in [Appendix A](#). The construct-validation work for individual features was performed by the tool authors: TAACO ([Crossley et al., 2016](#)) was built and validated for text cohesion, MRC ([Wilson, 1988](#)) for psycholinguistic norms (age of acquisition, concreteness, imageability), CEFR/EFLLex ([Dürlich and François, 2018](#)) for lexical sophistication, and so on. We aggregate these construct-validated features under Strohmaier's categorization rather than relabeling generic features.

3.3.1 Linguistic Simplicity

Lexical Simplicity reflects word-level difficulty during the construction stage, where readers decode individual words and terms. Unfamiliar words, technical jargon, and uncommon vocabulary are assumed to increase extraneous load at this stage, word frequency being associated with processing effort ([Berglund-Barraza et al., 2019](#)). We operationalize this dimension using psycholinguistic norms (age of acquisition, concreteness, imageability from MRC ([Wilson, 1988](#))), vocabulary sophistication (academic word coverage ([Cox-](#)

[head, 2000](#)), CEFR profiling ([Dürlich and François, 2018](#))), and word-level statistics (frequency, length, syllable counts).

Syntactic Simplicity captures sentence-level structural complexity during the construction and integration stages. Complex syntax (passive constructions, long dependency chains,...) increases extraneous load by requiring readers to maintain multiple open syntactic relations in working memory. We operationalize this using dependency tree properties (arc length, branching), tree depth ([Yngve, 1960](#)), clause counts, passive constructions, and sentence length measures.

Cohesion captures the explicit linguistic cues that help readers connect ideas during the integration stage, where propositions are linked into a coherent textbase. Things such as referential overlap, causal connectives, and semantic similarity reduce extraneous load by making relationships between sentences explicit ([McNamara and Kintsch, 1996](#)). We operationalize this primarily using TAACO features ([Crossley et al., 2016, 2019](#)): lexical overlap at sentence and paragraph level, connective counts, and givenness indicators from TAACO, and the semantic similarity indices (LSA, LDA, word2vec) added in TAACO 2.0.

Personalization and Motivation captures the use of direct address (*you, we*) and interpersonal language to make the text feel more personal and motivate the reader. Personalization works through a different route: conversational and personalized phrasing raises engagement and prompts deeper processing rather than changing a specific load type ([Mayer, 2014](#); [Cordova and Lepper, 1996](#); [Ginns et al., 2013](#); [Strohmaier et al., 2023](#)). We operationalize this using pronoun distributions (personal, possessive) and semantic category indicators for communicative, motivational, and affective language from the Harvard IV-4 dictionary ([Stone et al., 1966](#)). Our pronoun features cover the narrow operationalization of personalization in [Strohmaier et al. \(2023\)](#) (second-person personal and possessive pronouns), under which the meta-analytic effect was obtained.

3.3.2 Content Simplicity

Concept Simplicity captures the inherent difficulty of the ideas themselves, which influences the intrinsic load throughout all CI stages ([Sweller, 1988](#); [Sweller et al., 1998](#)). Regardless of how

Dimension	#Features		Primary sources
	Raw	Clean	
<i>Linguistic</i>			
Lexical	86	23	CEFR, AWL
Syntactic	64	18	Syntax Trees, TAACO
Cohesion	182	22	TAACO, Semantic Cat.
<i>Content</i>			
Coherence	47	11	TAACO, CoUDA, BBscore
Concept	42	7	CoCo, AWL
Person.	24	12	Semantic Cat.
Total	445	93	

Table 2: Dimension feature set summary.

the text is written, the cognitive effort required to read text depends on abstraction level, domain-specificity, and information density (number of concepts per unit of text). We operationalize it through entity metrics (entities per sentence, entity density, unique entity counts), concreteness and imageability norms from MRC (Wilson, 1988), and academic sophistication indicators from TextEvaluator (Sheehan et al., 2014) (nominalization rates, abstract noun ratios, academic word coverage).

Discourse Coherence captures the logical and semantic organization of ideas in the text. Unlike cohesion, which relies on explicit linguistic markers, coherence refers to the underlying conceptual structure. We look at whether ideas follow a logical progression (Kintsch, 1998). Strong coherence supports situation model construction by reducing germane load. We operationalize this using features derived from CoCo (Stajner et al., 2020), which was built to measure conceptual complexity through spreading activation over a knowledge graph. We measure the activation dynamics and connectivity signals of the resulting graph as indicators of how consistently a text develops its concepts. We combine these with Brownian Bridge coherence scoring (Sheng et al., 2024), CoUDA global and local coherence scoring (Zhu et al., 2024), and TAACO paragraph-level semantic similarity (LSA, LDA, word2vec) (Crossley et al., 2019).

3.4 Feature Processing

Many linguistic features are correlated with length, and word count correlates significantly with simplicity ratings. To isolate what dimension contribute beyond word length, we apply a normalization pipeline to each dataset independently. Precisely, we (1) residualize all features with respect to word count via OLS regression, removing the linear

effect of length; then (2) discard features with near-zero variance across all datasets; and (3) for each pair with high collinearity ($|r| > 0.85$), randomly remove one; finally, (4) we z-score all remaining features using the pooled source-simplified distribution per dataset. After cleaning and selection, 93 features remain (from 445 initial).

3.5 Validating Dimensions Against Ratings

To verify that our dimensions capture significant aspects of perceived simplicity, we examine their correlation with human ratings. For each dataset with simplicity ratings: we train Linear Regression and Random Forest regressors to predict ratings, using features from each simplicity dimension separately. We use 5-fold CV with randomized hyperparameter search and nested loops. We report the absolute Pearson correlation between predicted and actual ratings. This allows us to determine, for example, whether Lexical Simplicity predicts significantly better than Coherence on scientific text. To test whether dimensions differ significantly in predictive power within each dataset, we include Williams significance tests in appendix.

3.6 Effect Size of Simplification

For each feature, we compute the standardized mean difference $|\Delta z|$ between simplified and source texts, where features are z-scored per dataset using the pooled distribution. Averaging $|\Delta z|$ across features within each simplicity dimension yields a cross-dataset comparable measure of simplification magnitude per dimension. As validation, we train per-dimension binary classifiers (Random Forest) to discriminate source from simplified texts. Classification accuracy above random chance confirms that a dimension captures genuine information about transformations during simplifications.

3.7 Mapping Measures to Dimensions

Finally, we examine which dimensions existing evaluation measures capture. We compute scores from SARI, BLEU, ROUGE-L, FKGL, BERTScore, LENS-SALSA, and LLM-as-judge (Claude Haiku) and correlate measure scores with dimension scores across all texts in each dataset. Strong correlation indicates sensitivity to that aspect of simplicity. This analysis reveals measure biases and explains domain-specific performance: a measure capturing only lexical features will underperform when ratings are driven by coherence.

4 Results

4.1 RQ1: Correlation with Human Judgment

Table 3 shows the absolute Pearson correlation with human judgments of simplicity on the included datasets. We trained a Random Forest (RF) and a linear regression (LR) model on the entire set of features and on each group of features separately. We also trained stacked models by taking the dimension-level models as base learners, then training a meta-learner on their outputs. We provide baselines and established measures for reference, including SARI (Xu et al., 2016), D-SARI (Sun et al., 2021), BLEU (Papineni et al., 2002), ROUGE-L (Lin, 2004), learned reference-free measures (SLE (Cripwell et al., 2023), LENS-SALSA (Heineman et al., 2023)), and an LLM judge (Claude, whose configuration and prompt are given in Appendix I). Each measure is computed per sample and correlated with that sample’s rating: reference-based measures score against the dataset’s references, FKGL, SLE and LENS-SALSA score the simplification alone, and the Claude judge sees the source, the simplification and the references.

Our regressors are trained on the ratings, so we cannot fairly compare them to the fixed metrics. We include them for context, not as a fair baseline. On sentence-level datasets, the RF model trained on all features predicts human ratings strongly: on SDA, $r=0.74$ (vs. Claude’s 0.73); on ST_{sent} , 0.42 (vs. BLEU at 0.35); on D-Wiki, 0.60 (vs. Claude’s 0.52). However, on ST_{para} , all our models yield low correlations ($|r| \leq 0.13$), falling well below SARI (0.32). This suggests that the features, while predictive at the sentence level, do not transfer straightforwardly to paragraph-level evaluation on this dataset. Our models are trained in cross-validation per dataset, so learned weights need not transfer. A cross-dataset transfer analysis (Appendix E) suggests the dimension-level representations partially do: regressors trained out of domain, with a locally trained meta-learner, recover 80% of in-domain performance. Still, these results show that the selected features are salient enough to predict human ratings of simplicity (RQ1). On individual dimensions, taking *All Features* generally outperforms any single one. The only exception is ST_{para} where *Personalization / Motivation* (RF) is the only dimension exceeding all-features RF, though both correlations are very low. Stacking does not consistently improve over the base RF model, though

	Measure	SDA	ST_{sent}	ST_{para}	D-Wiki
Baselines	FKGL	0.07	0.32	0.01	0.09
	Word Count	0.03	0.34	0.02	0.34
	Avg word len.	0.07	0.00	0.17	0.06
	SLE	0.03	0.34	0.02	0.09
	BLEU	0.26	0.35	0.23	0.06
	ROUGE-L	0.28	0.33	0.15	0.21
	SARI	0.33	0.22	0.32*	0.37
	D-SARI	0.33	0.23	0.17	0.37
	Claude	0.73	0.18	0.19	0.52
RandomForest	All Features	0.74*	0.42	0.10	0.60*
	Lexical	0.73	0.37	0.03	0.55
	Syntactic	0.71	0.29	0.09	0.53
	Cohesion	0.21	0.32	0.05	0.38
	Concept	0.23	0.31	0.01	0.26
	Disc. Coherence	0.14	0.30	0.08	0.38
	Personalization	0.07	0.19	0.13	0.18
LinearRegression	All Features	0.56	0.02	0.12	0.27
	Lexical	0.60	0.16	0.08	0.42
	Syntactic	0.59	0.00	0.06	0.36
	Cohesion	0.19	0.00	0.07	0.12
	Concept	0.14	0.12	0.03	0.08
	Disc. Coherence	0.07	0.08	0.06	0.05
	Personalization	0.03	0.05	0.06	0.13
Stack	RF→RF	0.71	0.46*	0.18	0.59
	RF→LR	0.73	0.45	0.25	0.59
	LR→LR	0.60	0.14	0.12	0.40
	LR→RF	0.62	0.27	0.07	0.43

Table 3: Absolute Pearson correlation with human ratings. RF = RandomForest, LR = LinearRegression. **Bold**: best in category per dataset. *Best on dataset. Datasets whose units span more than one sentence are underlined.

RF→LR is competitive across datasets. The low performance on ST_{para} likely stems from its synthetic construction: paragraphs were formed by concatenating ST_{sent} sentences with averaged annotations (Davari et al., 2024), creating mismatch between paragraph-level feature computation and aggregated sentence-level ratings. Discourse features (cohesion, coherence) measure properties absent in ratings. The good performance of the D-Wiki (0.60) document-level dataset confirms this limitation is specific to the aggregation procedure, not to longer text spans.

4.2 RQ2: Analyzing Effect Size of Dimensions

Table 4 reports the average standardized mean difference ($|\Delta z|$) per dimension and dataset. Since features are z-scored per dataset using the pooled source–simplified distribution, $|\Delta z|$ represents the average shift in pooled standard deviations (SD).

Overall simplification intensity varies substantially across datasets, from 0.32 SD (SDA) and 0.40 SD (ASSET) to 0.82 SD (Cochrane). The

Dimension	ASSET	<u>OSE_{E→A}</u>	<u>OSE_{E→I}</u>	<u>OSE_{I→A}</u>	SDA	ST _{sent}	ST _{para}	<u>D-Wiki</u>	<u>Cochrane</u>	<u>Cochrane_{auto}</u>
FKGL (resid.)	0.80	0.89	0.70	0.72	0.65	0.91	0.63	0.70	0.88	0.74
Word Count	0.60	0.88	0.54	0.58	0.45	0.76	0.81	0.77	0.98	1.10
Lexical	0.54	0.93	0.66	0.64	0.46	0.56	0.61	0.69	0.90	0.76
Syntactic	0.44	0.74	0.56	0.53	0.37	0.46	0.64	0.62	0.90	0.77
Cohesion	0.37	0.75	0.57	0.53	0.27	0.40	0.59	0.56	0.90	0.78
Concept	0.46	0.82	0.61	0.57	0.42	0.50	0.58	0.62	0.88	0.75
Disc. Coherence	0.37	0.53	0.42	0.38	0.34	0.53	0.67	0.61	0.57	0.58
Personalization	0.31	0.73	0.55	0.51	0.25	0.33	0.44	0.47	0.69	0.59
Avg. All Features	0.40	0.74	0.56	0.52	0.32	0.43	0.58	0.58	0.82	0.71

Table 4: Average absolute effect size ($|\Delta z|$) per dimension and dataset. **Bold**: highest effect size dimension per dataset. Datasets whose units span more than one sentence are underlined. FKGL and Word Count included as baselines

OSE datasets show a clean gradient where the effect size scales with the proficiency gap between levels ($E \rightarrow A$: 0.74 > $E \rightarrow I$: 0.56 > $I \rightarrow A$: 0.52). Across datasets, *Lexical Simplicity* consistently shows the largest effect sizes, followed by *Syntactic Simplicity* in line with the traditional emphasis on these two dimensions in ATS research. *Concept Simplicity* follows closely, on par with or slightly below lexical and syntactic effects. Notably, *Personalization* shows the smallest effect size across most datasets, OSE being the exception. The effect size ranges from small (0.25, SDA) to moderate (0.59, Cochrane_{auto}), being high only on Cochrane (0.69) and OSE_{E→A}, despite being among the strongest predictors of human simplicity judgments in psychological research. This gap between scientifically-observed importance, and effect size suggests that current simplification practices under-target this dimension. *Cohesion* and *Discourse Coherence* similarly show below-average effect sizes, with the exception of the Cochrane dataset where Cohesion effects are large (0.90), equal with Lexical and Syntactic Simplicity; this may reflect the specific training and guidelines of the Cochrane PLS writers. Likewise, *discourse coherence* effect sizes scale with text span: they average 0.41 SD on the sentence-level datasets, 0.50 SD on the paragraph-level ones, and 0.59 SD on the document-level ones. Table 5 reports per-dimension classification accuracy (F1 macro) for discriminating source from simplified texts. Using all features, Random Forest separates source from simplified strongly on ST_{sent} (0.98) and the Cochrane datasets (0.91, 0.87), and moderately on the rest (0.61–0.86), with OSE hardest, consistent with their lower effect sizes (Table 4). ST_{para} is an exception: several per-dimension classifiers fall at




Figure 2: Mean absolute Pearson correlation ($|r|$) between per-dimension regressor predictions and established simplicity measures, averaged across the four datasets with per-sample human ratings (SDA, ST_{sent}, ST_{para}, D-Wiki).

or below chance (Cohesion 0.35, Personalization 0.36), and all-features RF reaches only 0.60 (LR all-features reaches 0.82). This is consistent with the RQ1 result and supports the same interpretation on text features and ratings mismatch. Setting ST_{para} aside, per-dimension classifiers confirm that each dimension carries simplification signal above baselines, with Personalization the weakest dimension on five of the nine remaining splits (RF F1 = 0.53–0.77), consistent with its small effect sizes (Table 4).

4.3 RQ3: measure correlation

Figure 2 shows the average correlation between measure scores and dimension-level predictions across four datasets. Correlations are low across the board: no traditional measure exceeds $|r| = 0.30$ on any dimension. More importantly, all measures show nearly flat profiles across dimensions. Among traditional measures, SARI shows the strongest correlations (0.22–0.30) but without meaningful differentiation between dimensions. Reference-based measures (BLEU, ROUGE-L,

	Measure	ASSET	<u>OSE_{E→A}</u>	<u>OSE_{E→I}</u>	<u>OSE_{I→A}</u>	SDA	ST _{sent}	ST _{para}	D-Wiki	Cochrane	Cochrane _{auto}
<i>Baselines</i>	Random	0.49	0.50	0.50	0.50	0.48	0.50	0.52	0.49	0.50	0.48
	Shuffle	0.50	0.51	0.51	0.51	0.53	0.51	0.52	0.53	0.50	0.52
	Avg Word Len.	0.59	0.63	0.52	0.55	0.61	0.59	0.70	0.62	0.53	0.54
	Word Count	0.58	0.62	0.57	0.56	0.58	0.54	0.47	0.65	0.68	0.72
	FKGL (resid.)	0.63	0.65	0.58	0.57	0.61	0.65	0.67	0.66	0.56	0.51
	FKGL	0.68	0.70	0.62	0.59	0.64	0.66	0.65	0.73	0.59	0.58
<i>RandomForest</i>	All Features	0.77*	0.78*	0.67*	0.61*	0.86	0.98	0.60	0.76*	0.91*	0.87*
	Lexical	0.56	0.70	0.61	0.58	0.58	0.57	0.64	0.65	0.73	0.68
	Syntactic	0.74	0.71	0.62	0.58	0.71	0.65	0.63	0.74	0.86	0.82
	Cohesion	0.65	0.70	0.60	0.58	0.61	0.61	0.35	0.68	0.81	0.79
	Concept	0.66	0.70	0.61	0.57	0.61	0.60	0.58	0.65	0.78	0.72
	Disc. Coherence	0.63	0.66	0.60	0.58	0.91*	0.98*	0.42	0.60	0.80	0.67
	Personalization	0.60	0.63	0.56	0.53	0.57	0.53	0.36	0.68	0.77	0.68
<i>LogisticRegression</i>	All Features	0.72	0.77*	0.67*	0.62*	0.71	0.97	0.82*	0.73	0.87	0.76
	Lexical	0.57	0.66	0.60	0.56	0.58	0.59	0.63	0.59	0.72	0.67
	Syntactic	0.69	0.66	0.60	0.57	0.68	0.63	0.76	0.71	0.76	0.68
	Cohesion	0.64	0.63	0.57	0.56	0.59	0.63	0.68	0.58	0.78	0.69
	Concept	0.61	0.67	0.60	0.57	0.60	0.57	0.65	0.62	0.75	0.69
	Disc. Coherence	0.61	0.60	0.57	0.53	0.55	0.97*	0.47	0.54	0.70	0.65
	Personalization	0.50	0.57	0.54	0.53	0.56	0.51	0.59	0.59	0.66	0.59
<i>Stack</i>	LR→RF	0.70	0.71	0.62	0.56	0.67	0.97	0.55	0.43	0.82	0.72
	RF→LR	0.76*	0.77	0.65	0.61	0.93*	0.98*	0.68	0.73	0.90	0.85

Table 5: F1 Macro classification scores. RF = RandomForest, LR = LogisticRegression. **Bold**: best in category per dataset. *Best on dataset. Datasets whose units span more than one sentence are underlined.

BERTScore) perform on par with length baselines, suggesting they respond to general text similarity rather than simplicity. Among neural approaches, SLE is the only measure showing a mild dimensional preference, with slightly higher sensitivity to Cohesion and Personalization. These flat profiles are notable given that dimensions vary substantially in their predictive power (RQ1) and transformation magnitude (RQ2). This means existing measures respond similarly to dimensions that matter most for comprehension and those that matter least. This limits their usefulness and may partly explain their poor correlation with simplicity ratings, which as shown in RQ1 are driven by different dimensions depending on the dataset. Because the six regressors use disjoint feature sets, nothing forces their predictions to agree. The flat profiles therefore reflect a lack of measure selectivity, not correlated dimensions. Unsupervised composites confirm this (Appendix G).

5 Conclusions

We showed that psycholinguistically-grounded simplicity dimensions can predict simplicity ratings, on most datasets, with dimension importance varying by population and text type (RQ1). We also showed that simplification datasets vary significantly in their transformation profiles: lexical and

syntactic dimensions dominate, while personalization (despite being a strong predictor of perceived simplicity) is consistently under-targeted (RQ2). Finally, we showed that existing measures capture little of this dimensional structure, correlating weakly and indiscriminately across all dimensions (RQ3). This work bridges psycholinguistic research on text comprehension with ATS evaluation. Grounding simplicity analysis in cognitive theory does more than motivate it: it predicts ratings, exposes blind spots current methods miss, and explains failures single-score measures cannot. We offer the framework as a basis for dimension-aware measures and better-balanced simplification datasets.

Limitations

There remains several limits to our claims, which we outline here

Circularity in Observed Dimensional Importance. While we argue (based on relevant psychology research) that dimensions beyond lexical and syntactic simplicity matter for comprehension, our empirical evidence for their importance is necessarily limited by the datasets available. The ATS field has historically focused on lexical and syntactic transformations, meaning existing datasets largely reflect this emphasis. As a result, dimen-

sions such as personalization, cohesion, and coherence exhibit limited variance in these datasets, which may attenuate both their observed effect sizes (RQ2) and their predictive power for human judgments (RQ1). The lower influence we observe for these dimensions may therefore reflect a data artifact from dataset design and implementation, rather than genuine unimportance. This circularity can only be perfectly resolved with datasets explicitly designed to vary along all dimensions. This is a field-level issue: Cardon and Doğruöz (2025) similarly documents weak and inconsistent correlations between readability measures, automatic metrics, and human judgments across ATS datasets, attributing this to the lack of a well-defined construct in the field.

Personalization Operationalization Scope. Our personalization features cover the narrow operationalization on which the meta-analytic effect was established (Strohmaier et al., 2023). Broader conceptions of personalization in discourse studies (engagement markers, hedges, boosters, and attitude markers in the sense of Hyland (2005)) are not captured by our lexical operationalization. Extending coverage to these would require dedicated NLP tools and dimension-specific evaluation data, and would also extend the construct beyond what the meta-analytic evidence supports. Although promising, the broader operationalization would carry weaker theoretical grounding than the current narrow one. Extending coverage in this direction remains an open research direction.

Evaluation of Under-Represented Dimensions. We also lack human-annotated datasets where raters were specifically trained to evaluate constructs such as cohesion or coherence in ATS. It is unclear whether human annotators can reliably assess these dimensions on a numerical scale, as opposed to holistic simplicity (annotators may not have a clear intuition of what an “8” or a “2” in cohesion means). Rather than relying on direct numerical assessment, one potential solution would be to evaluate along a distribution of texts known to span a wide difficulty range: from texts aimed at young children to scientific papers, through known popular outlets. Building such benchmarks remains an open challenge.

Feature and Dimensional Overlap. Several features load on multiple dimensions (marked with * in Table 6), meaning the dimensional separation is

not perfectly clean. While we train per-dimension models on disjoint feature *subsets* after assignment, the shared features introduce some conceptual overlap, notably between *Lexical* and *Concept* (e.g., both use features based on entity complexity). This limits the interpretability of per-dimension comparisons, but remains a reasonable trade-off given that these dimensions are not fundamentally independent constructs, and that our evaluators are not perfect. For example, entity complexity will inherently be correlated with the lexical complexity of the terms used to describe it, despite being seemingly independent conceptually.

Language Specificity. Our operationalization relies on English-specific resources: MRC psycholinguistic norms, CEFR profiling, the Academic Word List, and the Harvard IV-4 dictionary. The framework’s dimensional structure is mostly language-agnostic in principle, as it derives from universal cognitive models (CI and CLT), but the feature-level implementation does not transfer to other languages without substantial adaptation. Languages that differ significantly from English, such as logographic writing systems (e.g., Chinese) or languages with free word order (e.g., Turkish, Finnish), may require rethinking which features are appropriate for each dimension (especially syntax or cohesion).

Directionality of Changes. In RQ2, we measure the magnitude of change per dimension using $|\Delta z|$ but do not assess whether changes move *toward* simplicity. We assume that transformations in aligned datasets are predominantly simplifying, which is reasonable given their construction, but individual features may shift in the direction of greater complexity (e.g., increased nominalization despite overall simplification). Moreover, directionality is not always straightforward at the feature level: for instance, a decrease in `efllex_ratio_B1` may indicate simplification if compensated by an increase in `efllex_ratio_A1` (shifting vocabulary toward a lower CEFR level), but not if `efllex_ratio_C1` increases instead. A more fine-grained directional analysis would therefore require feature-level polarity annotations that account for such interactions, which we leave to future work.

Cognitive Mechanism Testing. The cognitive mechanisms behind CI and CLT (working memory load, propositional integration, expertise reversal,

and so on) can only be tested directly with methodology from psycholinguistic experimentation (eye-tracking, dual-task paradigms, fNIRS, and similar). Strohmaier et al. (2023) themselves do not test these mechanisms directly; their meta-analysis aggregates studies measuring comprehension outcomes after text modifications. Our framework operates at the same general level, using construct-validated text features grouped by these dimensions. Our results are about showing how predictive are these features grouped by these dimensions. We do not aim validate of the underlying cognitive theories.

Acknowledgments

This research was funded by the French National Research Agency (ANR) under the project ANR-22-CE23-0019-01. The authors acknowledge the use of *ChatGPT*, *Claude*, and *Grammarly* during the preparation of this work for the following purposes: assistance with code development, grammar and spelling correction, and rewording of text. All output generated with the aid of these tools was reviewed and revised by the authors, who bear full responsibility for the content of this publication.

References

- Sweta Agrawal and Marine Carpuat. 2024. [Do Text Simplification Systems Preserve Meaning? A Human Evaluation via Reading Comprehension](#). *Transactions of the Association for Computational Linguistics*, 12:432–448.
- Noof Abdullah Alfear, Dimitar Kazakov, and Hend Al-Khalifa. 2024. [Meta-Evaluation of Sentence Simplification Metrics](#). In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*, pages 11229–11235, Torino, Italia. ELRA and ICCL.
- Fernando Alva-Manchego, Louis Martin, Antoine Bordes, Carolina Scarton, Benoît Sagot, and Lucia Specia. 2020. [ASSET: A Dataset for Tuning and Evaluation of Sentence Simplification Models with Multiple Rewriting Transformations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 4668–4679, Online. Association for Computational Linguistics.
- Fernando Alva-Manchego, Carolina Scarton, and Lucia Specia. 2021. [The \(Un\)Suitability of Automatic Evaluation Metrics for Text Simplification](#). *Computational Linguistics*, 47(4):861–889.
- Jan Bakker and Jaap Kamps. 2024. [Cochrane-auto: An Aligned Dataset for the Simplification of Biomedical Abstracts](#). In *Proceedings of the Third Workshop on Text Simplification, Accessibility and Readability (TSAR 2024)*, pages 41–51, Miami, Florida, USA. Association for Computational Linguistics.
- Amy Berglund-Barraza, Fenghua Tian, Chandramalika Basak, and Julia L. Evans. 2019. [Word Frequency Is Associated With Cognitive Effort During Verbal Working Memory: A Functional Near Infrared Spectroscopy \(fNIRS\) Study](#). *Frontiers in Human Neuroscience*, 13:433.
- Rémi Cardon and A. Seza Doğruöz. 2025. [Readability Measures and Automatic Text Simplification: In the Search of a Construct](#). *Preprint*, arXiv:2511.09536.
- Diana I. Cordova and Mark R. Lepper. 1996. [Intrinsic motivation and the process of learning: Beneficial effects of contextualization, personalization, and choice](#). *Journal of Educational Psychology*, 88(4):715–730.
- Averil Coxhead. 2000. [A New Academic Word List](#). *TESOL Quarterly*, 34(2):213–238.
- Liam Cripwell, Joël Legrand, and Claire Gardent. 2023. [Simplicity Level Estimate \(SLE\): A Learned Reference-Less Metric for Sentence Simplification](#). In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 12053–12059, Singapore. Association for Computational Linguistics.
- Scott A. Crossley, Kristopher Kyle, and Mihai Dascalu. 2019. [The Tool for the Automatic Analysis of Cohesion 2.0: Integrating semantic similarity and text overlap](#). *Behavior Research Methods*, 51(1):14–27.
- Scott A. Crossley, Kristopher Kyle, and Danielle S. McNamara. 2016. [The tool for the automatic analysis of text cohesion \(TAACO\): Automatic assessment of local, global, and text cohesion](#). *Behavior Research Methods*, 48(4):1227–1237.
- Dennis Davari, Liana Ermakova, and Ralf Krestel. 2024. [Comparative analysis of evaluation measures for scientific text simplification](#). In *Linking Theory and Practice of Digital Libraries - 28th International Conference on Theory and Practice of Digital Libraries, TPDL 2024, Ljubljana, Slovenia, September 24-27, 2024, Proceedings, Part I*, volume 15177 of *Lecture Notes in Computer Science*, pages 76–91. Springer.
- Ashwin Devaraj, Iain Marshall, Byron Wallace, and Junyi Jessy Li. 2021. [Paragraph-level Simplification of Medical Texts](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4972–4984, Online. Association for Computational Linguistics.
- Luise Dürlich and Thomas François. 2018. [EFLLex: A Graded Lexical Resource for Learners of English as a Foreign Language](#). In *Proceedings of the Eleventh International Conference on Language Resources and*

- Evaluation (LREC 2018)*, Miyazaki, Japan. European Language Resources Association (ELRA).
- Liana Ermakova, Hosein Azarbondy, Jan Bakker, Benjamin Vendeville, and Jaap Kamps. 2026. [Overview of the CLEF 2025 SimpleText Track: Simplify Scientific Text \(and Nothing More\)](#). In Jorge Carrillo-de-Albornoz, Alba García Seco de Herrera, Julio Gonzalo, Laura Plaza, Josiane Mothe, Florina Piroi, Paolo Rosso, Damiano Spina, Guglielmo Faggioli, and Nicola Ferro, editors, *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, volume 16089 of *Lecture Notes in Computer Science*, pages 436–463. Springer Nature Switzerland, Cham.
- Liana Ermakova, Eric SanJuan, Stéphane Huet, Hosein Azarbondy, Giorgio Maria Di Nunzio, Federica Vezzani, Jennifer D’Souza, and Jaap Kamps. 2024. [Overview of the CLEF 2024 SimpleText Track: Improving Access to Scientific Texts for Everyone](#). In *Experimental IR Meets Multilinguality, Multimodality, and Interaction*, volume 14959 of *Lecture Notes in Computer Science*, pages 283–307, Cham. Springer Nature Switzerland.
- Rudolph Flesch. 1948. [A new readability yardstick](#). *Journal of Applied Psychology*, 32(3):221–233.
- D. Jake Follmer and Rayne A. Sperling. 2018. [Interactions between reader and text: Contributions of cognitive processes, strategy use, and text cohesion to comprehension of expository science text](#). *Learning and Individual Differences*, 67:177–187.
- Paul Ginns, Andrew J. Martin, and Herbert W. Marsh. 2013. [Designing Instructional Text in a Conversational Style: A Meta-analysis](#). *Educational Psychology Review*, 25(4):445–472.
- Arthur C. Graesser, Danielle S. McNamara, and Jonna M. Kulikowich. 2011. [Coh-Matrix: Providing Multilevel Analyses of Text Characteristics](#). *Educational Researcher*, 40(5):223–234.
- David Heineman, Yao Dou, Mounica Maddela, and Wei Xu. 2023. [Dancing Between Success and Failure: Edit-level Simplification Evaluation using SALSA](#). *Preprint*, arXiv:2305.14458.
- Ken Hyland. 2005. *Metadiscourse: Exploring Interaction in Writing*. Continuum, London and New York.
- J. Peter Kincaid, Robert P. Fishburne, Richard L. Rogers, and Brad S. Chissom. 1975. [Derivation Of New Readability Formulas \(Automated Readability Index, Fog Count And Flesch Reading Ease Formula\) For Navy Enlisted Personnel](#). Technical Report Research Branch Report 8-75, Naval Technical Training Command, Millington, TN. Research Branch.
- Walter Kintsch. 1998. *Comprehension: A Paradigm for Cognition*. Cambridge University Press, New York, NY, US.
- Christin Kreutz, Fabian Haak, Björn Engelmann, and Philipp Schaer. 2024. [BATS: Benchmarking Text Simplicity](#). In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 11968–11989, Bangkok, Thailand. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A Package for Automatic Evaluation of Summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Mounica Maddela, Yao Dou, David Heineman, and Wei Xu. 2023. [LENS: A Learnable Evaluation Metric for Text Simplification](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 16383–16408, Toronto, Canada. Association for Computational Linguistics.
- Richard E. Mayer. 2014. [Cognitive theory of multimedia learning](#). In Richard E. Mayer, editor, *The Cambridge Handbook of Multimedia Learning*, 2 edition, Cambridge Handbooks in Psychology, pages 43–71. Cambridge University Press, New York, NY, US.
- Danielle S. McNamara and Walter Kintsch. 1996. [Learning from texts: Effects of prior knowledge and text coherence](#). *Discourse Processes*, 22(3):247–288.
- Tenaha O’Reilly and Danielle S. McNamara. 2007. [Reversing the Reverse Cohesion Effect: Good Texts Can Be Better for Strategic, High-Knowledge Readers](#). *Discourse Processes*, 43(2):121–152.
- Yasuhiro Ozuru, Kyle Dempsey, and Danielle S. McNamara. 2009. [Prior knowledge, reading skill, and text cohesion in the comprehension of science texts](#). *Learning and Instruction*, 19(3):228–242.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. [Bleu: A Method for Automatic Evaluation of Machine Translation](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Günter Daniel Rey and Florian Buchwald. 2011. [The expertise reversal effect: Cognitive load and motivational explanations](#). *Journal of Experimental Psychology: Applied*, 17(1):33–48.
- Horacio Saggion. 2017. *Automatic Text Simplification*. Synthesis Lectures on Human Language Technologies. Morgan & Claypool Publishers.
- Kathleen M. Sheehan, Irene Kostin, Diane Napolitano, and Michael Flor. 2014. [The TextEvaluator Tool: Helping Teachers and Test Developers Select Texts for Use in Instruction and Assessment](#). *The Elementary School Journal*, 115(2):184–209.

- Zhecheng Sheng, Tianhao Zhang, Chen Jiang, and Dongyeop Kang. 2024. [BBScore: A Brownian Bridge Based Metric for Assessing Text Coherence](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(13):14937–14945.
- Sanja Stajner, Sergiu Nisioi, and Ioana Hulpus. 2020. [CoCo: A Tool for Automatically Assessing Conceptual Complexity of Texts](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 7179–7186, Marseille, France. European Language Resources Association.
- Philip J. Stone, Dexter Dunphy, Marshall S. Smith, and Daniel M. Ogilvie. 1966. *The General Inquirer: A Computer Approach to Content Analysis*. MIT Press, Cambridge, MA, USA. Harvard IV-4 word lists distributed via the General Inquirer repository.
- Anselm R. Strohmaier, Timo Ehmke, Hendrik Härtig, and Dominik Leiss. 2023. [On the role of linguistic features for comprehension and learning from STEM texts. A meta-analysis](#). *Educational Research Review*, 39:100533.
- Renliang Sun, Hanqi Jin, and Xiaojun Wan. 2021. [Document-Level Text Simplification: Dataset, Criteria and Baseline](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 7997–8013, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- John Sweller. 1988. [Cognitive Load During Problem Solving: Effects on Learning](#). *Cognitive Science*, 12(2):257–285.
- John Sweller, Jeroen J. G. van Merriënboer, and Fred G. W. C. Paas. 1998. [Cognitive Architecture and Instructional Design](#). *Educational Psychology Review*, 10(3):251–296.
- Leonard Tetzlaff, Bianca Simonsmeier, Tabea Peters, and Garvin Brod. 2025. [A cornerstone of adaptivity – A meta-analysis of the expertise reversal effect](#). *Learning and Instruction*, 98:102142.
- Sowmya Vajjala and Ivana Lučić. 2018. [On-eStopEnglish corpus: A new corpus for automatic readability assessment and text simplification](#). In *Proceedings of the Thirteenth Workshop on Innovative Use of NLP for Building Educational Applications*, pages 297–304, New Orleans, Louisiana. Association for Computational Linguistics.
- Benjamin Vendeville, Liana Ermakova, and Pierre De Loor. 2025. [Resource for Error Analysis in Text Simplification: New Taxonomy and Test Collection](#). In *Proceedings of the 48th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR ’25*, pages 3723–3732, New York, NY, USA. Association for Computing Machinery.
- Michael Wilson. 1988. [MRC psycholinguistic database: Machine-usable dictionary, version 2.00](#). *Behavior Research Methods, Instruments, & Computers*, 20(1):6–10.
- Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. 2016. [Optimizing Statistical Machine Translation for Text Simplification](#). *Transactions of the Association for Computational Linguistics*, 4:401–415.
- Victor H. Yngve. 1960. [A model and an hypothesis for language structure](#). *Proceedings of the American Philosophical Society*, 104(5):444–466.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [BERTScore: Evaluating Text Generation with BERT](#). In *8th International Conference on Learning Representations (ICLR 2020)*, Addis Ababa, Ethiopia. OpenReview.net.
- Dawei Zhu, Wenhao Wu, Yifan Song, Fangwei Zhu, Ziqiang Cao, and Sujian Li. 2024. [CoUDA: Coherence Evaluation via Unified Data Augmentation](#). 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 967–978, Mexico City, Mexico. Association for Computational Linguistics.

A Feature List per Dimension

B Framework Comparison

Table 7 summarizes the comparison. The frameworks differ less in feature coverage than in how features are sourced and dimensions are organized. Our framework aggregates construct-validated features (TAACO for cohesion, MRC for psycholinguistic norms, CEFR/EFLLex for lexical sophistication, etc.) under dimensions derived deductively from CI/CLT theory. TextEvaluator extracts eight components via PCA on a heterogeneous feature pool and labels them post-hoc; BATS aggregates 37 binary features from the literature without organizing them into dimensions. TextEvaluator additionally covers narrativity and argumentation, which fall outside the simplicity-focused scope of our framework.

C Length Residualization

C.1 Motivation

Many linguistic features are inherently correlated with text length (e.g., longer texts naturally have more clauses, more discourse markers, higher type-token ratios at certain scales). If left uncorrected,

Table 6: Features retained after preprocessing, grouped by dimension. Features marked with * load on multiple dimensions.

Dimension	Features	#
Lexical Simplicity	avg_age_of_acquisition, dale_chall, smog, total_syllables eflex_ratio_{a1, a2, b1, b2, c1} proportion_long_words_{letters_9, syllables_3, syllables_4} ratio_infrequent_words_bats taaco_{adj, adv, bigram_lemma, lemma, noun, prp, verb}_ttr teawl_ratio*, te_ratio_nominalizations_token* unieval_fluency*	23
Syntactic Simplicity	avg_words_before_main_verb, per_word_noun_phrases, token_count ratio_{appositions, conditionals, modifiers, negations} ratio_passive_voice_ext, ratio_past_tense_ext* ratio_past_perfect_{bats, ext} ratio_relative_clauses_{bats, ext} ratio_subordinate_clauses te_avg_{dependent_clauses, yngve_depth}, te_freq_long_sentences_25w unieval_fluency*	18
Cohesion	avg_distance_same_entities*, coco_avgTokenDistance* taaco_addition, taaco_opposition, taaco_order, taaco_sentence_linking taaco_adjacent_overlap_pronoun_sent taaco_all_{causal, demonstratives, logical, negative, positive, temporal} taaco_{coordinating_conjuncts, determiners, lexical_subordinators} taaco_lsa_1_all_sent taaco_positive_{causal, intentional, logical} te_ratio_causal_conjuncts, unique_entity_total_entity_ratio_text	22
Discourse Coherence	avg_distance_same_entities*, bbsscore, couda_global, paragraph_count coco_{activationDecayRate, activationMax, activationMin, activationStability} coco_avgNodeDegree, coco_avgTokenDistance*, coco_complexityScore*	11
Concept Simplicity	coco_complexityScore*, coco_entityDensity, entities_per_sentence ratio_{cog_orient, institution}_gi teawl_ratio*, te_ratio_nominalizations_token*	7
Personalization & Engagement	ratio_{active, communicating, motivation, negativ, our}_gi ratio_{ovrst, passive, positiv, self, undrst, you}_gi te_ratio_pronouns_3rd_singular*	12
Total unique features		93

a model trained on these features may learn to predict simplicity primarily through length rather than through genuine linguistic simplification properties. We use length residualization so we can study the effect of the variance not attributable to text length, which can easily be studied separately.

C.2 Procedure

The procedure is as follows:

1. **Identifying length-correlated features:** For each feature f in the feature matrix, the Pearson correlation r between f and the length proxy (word count) is computed. A feature is marked for residualization if $|r| > 0.1$ (a deliberately low threshold to be conservative so even weak length confounds are removed).
2. **Fitting an OLS regression:** For every feature f flagged in Step 1, an Ordinary Least-

Squares model is fitted:

$$f = \beta_0 + \beta_1 \cdot \text{word_count} + \varepsilon$$

where β_0 is the intercept, β_1 is the regression coefficient capturing the linear relationship between the feature and word count, and ε represents the residuals. We use OLS via the ‘statsmodels.OLS’ library, with an intercept term.

3. **Replacing the feature with the residuals:** The original feature values are replaced by the OLS residuals ε . These residuals represent the portion of the feature’s variance that is *orthogonal to text length* i.e., the feature value after the linear effect of length has been subtracted.

C.3 Grouping Mode

Residualization is performed per-dataset: for paired datasets (e.g., source and generated texts

	Ours	TextEvaluator	BATS
Lexical	✓	C1, C4	✓
Syntactic	✓	C2	✓
Cohesion	✓	C7	partial
Concept Simplicity	✓	C3 (partial)	partial
Discourse Coherence	✓	–	–
Personalization	✓	C5, C6 (partial)	–
Dimensions not in ours	–	C6 (narrativity), C8 (argumentation)	–
Feature provenance	construct-validated tools	feature pool	literature aggregation
Dimension construction	deductive from theory	PCA on features	–
Theoretical grounding	CI/CLT comprehension theory	–	–

Table 7: Comparison across multidimensional ATS evaluation frameworks. TextEvaluator components labelled C1–C8 following the component order in Sheehan et al. (2014). “Partial” indicates partial coverage of the dimension.

from the same corpus), both roles are pooled together before fitting the OLS model. This is important because fitting separate regressions per role (source vs. generated) would introduce artificially different intercepts/slopes between the two groups, creating spurious separability in downstream discrimination analyses. By using a single shared regression across both roles of a dataset, genuine linguistic differences between source and simplified texts are preserved.

C.4 Effect

After residualization, the mean absolute correlation between residualized features and word count drops to approximately zero, confirming that the linear length confound has been effectively removed. Features that were not correlated with length ($|r| \leq 0.1$) are left unchanged.

D Williams Significance Tests

Figure 3 reports pairwise Williams significance tests comparing the Pearson correlations of all models with human judgments, for each annotated dataset. A green cell indicates that the row model correlates significantly more than the column model ($p < 0.05$). Models are sorted by descending $|r|$.

Several patterns emerge. First, on SDA and D-Wiki, the top-performing models (All Features RF, Lexical RF, Syntactic RF, and stacked variants) form a cluster with no significant pairwise differences, indicating that Lexical and Syntactic features alone capture most of the predictive signal on those datasets. On STsent, by contrast, stacked models significantly outperform individual dimensions, suggesting that multidimensional combination is most beneficial when no single dimension

dominates human judgments. Second, RF significantly outperforms LR on the same feature sets across most datasets, indicating that nonlinear feature interactions contribute meaningfully to prediction — a relevant consideration for future measure design.

Third, Personalization / Style ranks significantly below most other dimensions on SDA, ST_{sent}, and D-Wiki. Combined with the small effect sizes reported in RQ2 (Table 4), this suggests that current datasets not only under-target personalization but that the resulting feature variance is too limited to reliably predict human ratings. Finally, the relative ordering of dimensions shifts across datasets: Cohesion and Discourse Coherence rank significantly below Lexical Simplicity on SDA but are not significantly different from it on ST_{sent}. This confirms that simplicity is not a fixed hierarchy of dimensions but a dataset-dependent structure, as discussed in Section 4.1.

E Out-Of-Domain Generalization

Dataset	Out-Of-Domain			In Domain	
	AF	Stack	Hybrid	AF	Stack
SDA	0.61	0.57	0.64	0.74	0.74
ST-sent	0.12	0.11	0.57	0.74	0.75
ST-para	0.29	0.30	0.34	0.35	0.40
D-Wiki	0.43	0.45	0.50	0.71	0.69
Mean	0.34	0.36	0.51	0.64	0.64

Table 8: Transfer strategy results by dataset. Best single donor as foreign source. AF being “All Features”

We want to evaluate whether dimension-level representations generalize across datasets. The question here is whether the per-dimension regressors learn transferable properties of simplicity, or




Figure 3: Pairwise Williams significance tests. A green cell indicates the row model correlates significantly more than the column model ($p < 0.05$) with human judgments.




Figure 4: Cross-dataset transfer performance (|Pearson r |) on different test dataset. Multiple setups are used including completely foreign training (Foreign All-Features, full foreign stacking), full in-domain (In-Domain Stacking, In-Domain All-Features) and Hybrid Stacking, where dimension classifiers trained Out-Of-Domain, are used to train an in-Domain Stacked regressor

just overfit to dataset-specific annotation patterns. To study this, we compare five transfer strategies with increasing access to target-domain data. Foreign All-Features and Full Foreign transfer an entire model trained on a foreign dataset, either as a single regressor or as a full stacking pipeline. These establish a lower bound on transferability. Hybrid is the critical condition: it retains the foreign dimension-level base regressors but retrains only the meta-learner on the target dataset via cross-validation. If dimensions capture generalizable signal, then recalibrating their combination locally should recover most of the in-domain performance, even when the base regressors have never seen the target data. Results are shown in Figure 4 and Table 8. This is what we observe: Hybrid recovers 80% of in-domain performance (0.51 vs 0.64), while full foreign transfer remains weak (0.36). The gap confirms that dimension-level representations transfer across datasets but their relative importance does not, consistent with the dataset-dependent patterns observed in RQ1.

F Pairwise Dimension Correlation

Figure 5 reports pairwise Pearson correlations between per-dimension regressor predictions on the four rated datasets, and Table 9 reports the cross-dataset average. Profiles vary substantially across datasets, reflecting that simplification operations are dataset-specific rather than dimension-specific: a dataset that consistently co-modifies two dimensions will produce correlated dimension scores even if the underlying constructs are separable. On SDA and D-Wiki, correlations are modest ($\max(|r|) = 0.19$ and 0.34 respectively). ST_{sent} is mostly modest but with one notable exception: Discourse Coherence and Concept Simplicity correlate at $|r| = 0.63$, likely reflecting features that load on both dimensions (marked with * in Table 6). ST_{para} shows substantially higher correlations across most pairs ($\max(|r|) = 0.71$ between Cohesion and Syntactic Simplicity), consistent with the structural problems with this dataset identified in §4.1 and §4.2. Averaged across datasets, correlations are modest ($\max(|r|) = 0.31$), with Lexical Simplicity showing the lowest average overlap with other dimensions.

An important caveat: simplification datasets produce some correlation between dimension scores even for perfectly orthogonal constructs, because simplifications often target multiple dimensions si-

	Lex	Syn	Coh	Con	Disc	Pers
Lexical	1.00	-	-	-	-	-
Syntactic	0.21	1.00	-	-	-	-
Cohesion	0.24	0.31	1.00	-	-	-
Concept	0.15	0.22	0.23	1.00	-	-
Disc. Coherence	0.16	0.18	0.31	0.30	1.00	-
Personalization	0.09	0.16	0.17	0.18	0.15	1.00

Table 9: Cross-dataset averaged correlations between dimension regressor predictions.

multaneously (a text receiving a lexical edit usually also receives a syntactic edit). The observed correlations therefore reflect correlated transformations in the data, not necessarily correlated constructs. Separating these would require datasets designed to vary dimensions independently, which we discuss in Limitations. The matrices should be read as an upper bound on construct-level overlap rather than a direct measurement of it.

G Composite Scores Correlation with Metrics

In the main analysis (Section 4.3), dimension scores are obtained from per-dimension regressors trained on human judgments. While these regressors are trained on disjoint feature subsets, meaning differential correlations across dimensions may reflect genuine measure sensitivity rather than just shared prediction targets. We verify this by replicating the analysis using two unsupervised alternatives that involve no training on human ratings.

The first method aggregates features within each dimension by averaging their z-scores across all features belonging to that dimension, for each text independently. This produces a single unsupervised score per dimension per text, where all features contribute equally. (Figure 7).

The second extracts the first principal component per dimension, capturing the dominant axis of variance within each feature group (Figure 6).

Absolute correlations are lower in both cases ($\max(|r|) \approx 0.13$ for z-composite, ≈ 0.18 for PCA, compared to ≈ 0.38 for regressor-based scores), which is expected: unsupervised aggregation weights all features equally or by variance rather than by relevance to perceived simplicity, diluting signal from the most informative features. Despite this, the key patterns from the main analysis hold. First, Lexical Simplicity remains the dimension most strongly associated with the majority of measures, including FKGL, SLE, and Claude-




Figure 5: Pairwise Pearson correlations between dimension regressor predictions on the four rated datasets.




Figure 6: Mean absolute Pearson correlation (lrl) between measures and unsupervised dimension scores (first PCA component per dimension), averaged across 4 datasets (D-Wiki, SDA, ST-para, ST-sent).




Figure 7: Mean absolute Pearson correlation (lrl) between measures and unsupervised dimension scores (mean of z-scored features per dimension), averaged across 4 datasets (D-Wiki, SDA, ST-para, ST-sent).

based evaluators. Second, no measure shows strong selective sensitivity to any single dimension — profiles remain relatively flat across dimensions. Third, Discourse Coherence, Concept Simplicity, and Personalization remain poorly captured by most evaluation measures. The convergence of supervised and unsupervised approaches on these patterns strengthens our conclusion that existing measures are weakly and indiscriminately coupled to the dimensional structure of simplicity.

H Reproducibility

All measures reported in this paper are implemented in a unified Python library, SIMPSY, released at <https://github.com/bVendeville/Simpsy>. It covers the 445-feature pool described in §3.3, the 93-feature residualized subset after preprocessing, and the baseline measures (FKGL, BLEU, SARI, D-SARI, BERTScore, LENS-SALSA, SLE). The analysis code and data are released separately at <https://github.com/bVendeville/beyond-lexical-and-syntactic>, comprising the dimension regressors and classifiers used in §4.1 and §4.2, the evaluation notebooks for reproducing the analyses, and the scored datasets for the four rated and six aligned datasets used in this paper.

I LLM Evaluation

We use Claude Haiku (claude-haiku-4-5-20251001) as our LLM-as-a-judge. Table 10 summarizes the configuration.

Parameter	Value
Model	claude-haiku-4-5-20251001
Temperature	0
max_tokens	10
Rate limiting	20 requests/second
Max retries	3 (exponential backoff)

Table 10: Details of the claude LLM-as-a-judge evaluation setup

Listing 1: Claude Prompt

Given the original text, its simplification,
 → and human reference simplifications,
 → rate the simplification on a scale of
 → 1-10:

- 1 = Poor (doesn't improve readability, loses
→ meaning, or much worse than references)
 - 10 = Excellent (greatly improves readability,
→ preserves meaning, comparable to
→ references)
- Original: "{orig}"
Simplified: "{simp}"
Human references: "{refs}"
Respond with ONLY a single integer from 1 to 10.

J Datasets Samples

SDA - Score: 37.87/100, higher=simpler

Source: Gable also earned an Academy Award nomination when he portrayed Fletcher Christian in 1935's *Mutiny on the Bounty*.

Simplified: Gable earned an Academy Award nomination he portrayed Fletcher Christian.

ASSET

Source: Prior to the arrival of the storm, the National Park service closed visitor centers and campgrounds along the outer banks.

Simplified: Before the storm came to the National Park service, the National Park service closed places and campgrounds along the outer banks.

ST-sent - Score: 1/14
syntax=1/7, lexical=1/7
lower=simpler

Source: Inappropriate use of medicines increases the risk of hospital admissions for the elderly.

Simplified: misusing medicines harms the elderly.

ST-para - Score: 2.165/14
syntax=1/7, lexical=2.33/7
higher=simpler

Source: As various kinds of output devices emerged, such as highresolution printers or a display of PDA(Personal Digital Assistant), the importance of high-quality resolution conversion has been increasing. This paper proposes a new method for enlarging image with high quality. One of the largest problems on image enlargement is the exaggeration of the jaggy edges. To remedy this problem, we propose a new interpolation method, which uses artificial neural network to determine the optimal values of interpolated pixels. The experimental results are shown and evaluated. The effectiveness of our methods is discussed by comparing with the conventional methods.

Simplified: in this paper, we propose a new method for resolution conversion paper proposes a new method for enlarging image with high quality. image enlargement's exaggeration of jaggy edges is one of the largest problems. to fix this issue, we suggest a new method that uses a computer brain to find the best values for filling in pixels. experimental results are shown and evaluated. efecundity of our methods is discussed by comparing with the conventional methods.

D-Wiki - Score: 4/5, higher=simpler

Source: on 22 january 2019 , floods caused by heavy rainfall struck the indonesian province of south sulawesi . at least 68 people were killed and thousands were displaced . gowa regency was particularly hard-hit , experiencing the majority of the fatalities.the most destructive flood was caused by the overflowing of the jeneberang river and the opening of the watergates of the bili-bili dam in gowa , though other floods also occurred elsewhere in the province .

Simplified: on 22 january 2019 , floods caused by heavy rainfall hit the indonesian province of south sulawesi . at least 68 people were killed and thousands were lost . gowa regency was hard-hit having the majority of the fatalities.the most destructive flood was caused by the overflowing of the jeneberang river and the opening of the watergates of the bili-bili dam in gowa , though other floods also occurred elsewhere in the province .