Mapping Personas to Text Transformations: A Taxonomy Outline for Content Adaptation

Andrea Sillano
Politecnico di Torino
Torino, TO, Italy
andrea.sillano@polito.it

Luigi De Russis
Politecnico di Torino
Torino, TO, Italy
luigi.derussis@polito.it

Tommaso Calò
Politecnico di Torino
Torino, TO, Italy
tommaso.calo@polito.it

Raphael Troncy
EURECOM
Sophia Antipolis, France
raphael.troncy@eurecom.fr

Pasquale Lisena
EURECOM
Sophia Antipolis, France
pasquale.lisena@eurecom.fr

Abstract

Tailoring writings to specific audiences increases engagement and comprehension, and LLMs with persona-based adaptation can make this process easier and scalable. However, current systems treat personas as a black-box prompt modifier: non-expert users can specify "simplify for a beginner" or "adapt for an expert", but cannot inspect how and what persona traits map to concrete textual edits. This opacity limits control, and prevents systematic understanding of what makes adaptations effective. We argue for decomposing persona-driven adaptation into transparent edit operations. Rather than monolithic rewrites, we outline a taxonomy that explicitly maps persona dimensions to auditable operations. Each operation becomes a traceable composition of operators, which specify personas, applied transformations, and design rationale, making adaptation choices visible and controllable.

CCS Concepts

• **Human-centered computing → Human computer interaction (HCI)**; • **Computing methodologies → Natural language processing**.

ACM Reference Format:
Andrea Sillano, Luigi De Russis, Tommaso Calò, Raphael Troncy, and Pasquale Lisena. 2026. Mapping Personas to Text Transformations: A Taxonomy Outline for Content Adaptation. In *Proceedings of Responsible Use of AI Personas in Human-Centered Design and Research Workshop of the 2026 CHI Conference on Human Factors in Computing Systems (CHI Workshop '26).* ACM, New York, NY, USA, 4 pages.

1 Introduction

Adapting writing to a particular audience makes material more engaging and easier to understand, but doing it manually is labor-intensive. Works in educational practice emphasize that comprehension improves when instruction explicitly supports readers' processing rather than assuming a one-size-fits-all text will work for all [1, 6]. With the advent of Large Language Models (LLMs), content adaptation has become easier and more scalable: systems can rewrite explanations, adjust reading level, and contextualize materials for a reader's background or interests. To make these edits possible, the model rewrites the passage acting like a target persona it is instructed to follow. LLMs enable the generation of multiple personas that can be used across different tasks [18]. Recent work explicitly targets this promise, for example, frameworks that personalize materials to a student's context aiming to reduce disengagement from standardized content [9]. In parallel, other work shows how persona prompting and role prompting with a defined audience profile can steer LLMs' output [16]. However, empirical evidence also suggests personas can change outputs in unpredictable ways and do not reliably improve performance, raising questions about when persona conditioning is beneficial, and what mechanisms it actually induces [18, 21]. In many current workflows, personas function as a black-box control: users must guess how a persona specification will be interpreted and which aspects of a text will change. The link between persona traits and concrete output transformations is rarely explicit or auditable. Persona conditioning is not neutral, as they can induce systematic biases, amplify stereotypes, and behave inconsistently across prompt formulations [7, 19]. As a result, inspecting how persona traits map to concrete textual edits is a non-trivial task. This opacity limits the ability to control which aspects of the text change (e.g., terminology, structure, evidence, cultural references). We argue that progress in trustworthy adaptation requires decomposing persona-driven rewrites into transparent, named transformation operators. Instead of monolithic "*Acts as a <persona description> rewrite this for <settings>*", we outline a literature informed-taxonomy that maps persona dimensions to auditable operations. This operator framing matches work on controllable simplification, which treats rewriting as a set of controllable knobs [11, 12, 20]. By making adaptation choices visible and auditable through the proposed mapping, we aim to shift persona-based personalization from an opaque prompt trick into a controllable and evaluable design space, opening possibilities for writers-configurable adaptation pipelines, thus improving transparency while keeping agency in the process.

2 Method and Taxonomy

To address the issue of black-box persona prompting for text transformation, we outline a literature-informed taxonomy that makes persona-based adaptation explicit by linking persona dimensions to

CHI Workshop '26, Barcelona, Spain

Table 1: Taxonomy of Persona Dimensions to Text Transformation Operations

Persona Dimension	Text Transformation Operation	Result	Example
Demographic Attributes	Lexical Structural Style	Vocabulary sophistication Sentence complexity adjustment Register formality	"utilize" → "use" (for general audience) Split compound sentences for novice readers Adjust professional vs. casual language
Domain Knowledge	Lexical Content	Technical terminology density Domain-specific elaboration	"Myocardial infarction" ↔ "heart attack" Add/remove technical explanations and definitions
Cognitive Attributes	Structural Content Tone	Syntactic complexity control Information granularity Instructional explicitness	Adjust clause embedding and dependency depth Vary abstraction level and detail density Add guidance ("Note that...", "This means...")
Contextual Attributes	Lexical Content	Cultural-linguistic localization Cultural reference adaptation	Adapt idioms, units, spelling conventions Replace culture-specific examples with familiar ones

specific text edits rather than implicit transformation. To propose such a taxonomy, we draw on two established literature research on persona modeling in LLMs and controllable text operations.

Persona Dimension. We ground our persona traits in recent works that decompose personas as structure, multi-dimensional representation [3, 17, 19]. We synthesize a layered persona model covering: (1) *demographic attributes*, (2) *domain knowledge*, (3) *cognitive attributes*, and (4) *contextual attributes* [3, 4]. These four dimensions were selected through a convergent synthesis as they represent the axes that most consistently showed across persona modeling frameworks and most reliably tends to produce systematic variance in LLM outputs [3, 5, 15]. We deliberately scope the taxonomy to dimensions with textual correlation, that is, dimensions for which a corresponding transformation operation can be concretely defined. Other attributes such as persona background and beliefs, while present in some persona models, are excluded at this stage because their mapping to discrete text operations remains underspecified. Empirical observations suggest that persona prompting is sensitive to prompt formulation, and systematic evaluation favors designs that isolate persona factors. Accordingly, structuring personas into separable dimensions can improve controllability and support clearer attribution compared to monolithic persona descriptions [10, 19].

Text Transformation Operation. We derived transformation types from controllable text literature, which has categorized rewriting operations at different granularity, identifying core operations across rewriting tasks: *lexical operation*, *structural operation*, and *content operation* [2, 11, 12, 20]. Furthermore, we move beyond text simplification and ground our operator on LLM style control to capture shifts in other attributes (e.g., tone, stance, formality, details) [8, 14]. This expanded view recognizes that persona-driven adaptation requires text-level operations, style modulation like, and tone adjustments.

By synthesizing recurring aspects from the literature described in the paragraphs above, we propose mappings between persona dimensions and a set of transformation operations. For instance, low expertise level triggers terminology simplification and sentence splitting, while domain background determines whether technical terms require definition or replacement. Cultural context influences example selection and cultural reference adaptation. Beyond these traditional simplification mappings, persona attributes also drive style transformations: professional role influences formality and tone, reading goals determine level of detail, and cognitive preferences shape sentence complexity and discourse markers. This mapping process builds on evidence that persona attributes produce variance in outputs [5, 15]: demographic and attitudinal factors systematically correlate with linguistic choices, suggesting that persona-to-text transformations follow predictable patterns rather than arbitrary mappings. Moreover, when adaptation takes place via a monolithic rewrite (e.g., "*Rewrite this passage as an experienced comedy writer*"), detecting if a bias is introduced (e.g. a cultural reference substitution that stereotypes a group) requires comparing input and output holistically, with neither intermediate control signal nor user agency. Under the operator framing, each transformation happens as a discrete, named step with an associated persona trigger and design rationale. A writer can therefore isolate the adaptation operator made by a given persona dimension, and evaluate it independently of unrelated changes to sentence complexity or formality. To make the mapping explicit, Table 1 presents a first outline of the mapping taxonomy that links selected persona dimensions to the set of text transformation operations. The table organizes operations into lexical, structural, content, style, and tone categories, and illustrates each mapping with a representative example. This taxonomy serves as a design scaffold for selecting transformation operators conditioned on persona attributes, aiming to reduce the intuition needed with persona-driven rewriting by making the process more systematic and interpretable edit patterns.

To illustrate how operators compose in practice, consider adapting the sentence *"The second law of thermodynamics states that the total entropy of an isolated system increases monotonically over time, with irreversible processes dissipating free energy and driving the system toward thermodynamic equilibrium."* with goal of making the sentence accessible to a general audience while retaining scientific credibility for a semi-expert reader. Consulting Table 1, it is possible to assemble a targeted operator chain by selecting individual operators from the two persona profiles defined in Table 2 and Table 3, each chosen because it can best serves the adaptation

Table 2: P1 - Persona profile

Dimension	Attribute
Demographic	High school student, age 16
Domain knowledge	No prior exposure to physics or thermodynamics
Cognitive	Novice reader; limited tolerance for syntactic complexity; requires concrete anchoring
Contextual	Practical reading goal: understand a concept for a class assignment
Op. contribution	Lexical simplification, Structural splitting, Inferential scaffolding

Table 3: P2 - Persona profile

Dimension	Attribute
Demographic	Science journalist, early career
Domain knowledge	General scientific literacy; familiar with formal register but not domain specialist
Cognitive	Comfortable with complex syntax; values evidence and precision
Contextual	Instrumental reading goal: accurately report a concept to a broad audience
Op. contribution	Content precision retention, register tightening

goal at that step. The result is a five-step composition where every textual change is traceable to a named operator and source persona.

Each step below was produced by prompting GPT-5.4 [13] to adopt the persona specified in the step header, the model was instructed to rewrite the text applying only the operator indicated and to leave all other aspects of the text unchanged.

Original. **The second law of thermodynamics**[Orig] states that the **total entropy**[Orig] of an **isolated system**[Orig] increases **monotonically**[Orig] over time, with **irreversible processes dissipating free energy**[Orig] and driving the system toward **thermodynamic equilibrium**[Orig].

*Step 1 — **Lexical**[Lex] ← P1 (Domain knowledge: none). Replace technical terms with plain equivalents.*
The second law of thermodynamics states that the total **disorder**[Lex] of a **closed system**[Lex] **always**[Lex] increases over time, with processes that cannot be undone **releasing stored energy**[Lex] and driving the system toward a **stable state**[Lex].

*Step 2 — **Content**[Cont] ← P2 (Domain knowledge: expert). Restore technically precise term where accuracy matters.*
The second law of thermodynamics states that the total disorder of a closed system always increases over time, with processes that cannot be undone **dissipating free energy**[Cont] and driving the system toward a stable state.

*Step 3 — **Structural**[Str] ← P1 (Cognitive: novice). Split compound sentence at natural clause boundary.*
The second law of thermodynamics states that the total disorder of a closed system always increases over time. **Irreversible processes dissipate free energy,**[Str] driving the system toward a stable state.

*Step 4 — **Style**[Sty] ← P2 (Demographic: professional). Tighten register, restore precise nominal form.*
The second law of thermodynamics states that the total disorder of a closed system always increases over time. Irreversible processes dissipate free energy, **driving the system toward a stable equilibrium state**[Sty].

*Step 5 — **Tone**[Tone] ← P1 (Contextual: practical goal). Add inferential scaffold anchored in everyday intuition.*
The second law of thermodynamics states that the total disorder of a closed system always increases over time**, in other words, things naturally tend toward chaos rather than order**[Tone]. Irreversible processes dissipate free energy, driving the system toward a stable equilibrium state.

Legend: Lex Lexical Cont Content Str Structural Tone Tone Sty Style

3 Conclusion and future work

This paper brings to light the issue of opaque persona-based content adaptation and its relevance for both system designers and end users. We aim to push for a shift towards a more informed use of personas by proposing a transition from black-box prompting to more transparent transformations. By mapping relevant persona dimensions to specific text operations, we outline a taxonomy that makes adaptation choices more visible and understandable. This work serves as an initial framework to understand how persona attributes can systematically correlate with textual change. Making operators visible can improve understanding over operation performed onto text, directly impacting human agency. This framework makes users more proactive since they are informed about *what* is being changed, *why* and under certain level *how*. Instead of passively accepting or rejecting a monolithic rewrite, users equipped with explicit operator mappings can intervene at the level of individual transformation choice improving agency over written text edits. We see this as a starting point for moving persona-based adaptation toward more systematic and auditable practices, and can foresee how this can have an impact on writing support tools, educational environment or user personalization. However, significant work remains to validate these mappings empirically and refine the taxonomy through real-world applications.

For this reason, we invite the community to: (1) Refine and extend this taxonomy; (2) Conduct empirical studies testing these persona-to-operation mappings with real users and diverse content types; (3) Develop evaluation frameworks for assessing adaptation validity.

References

[1] Leah L. Anderson, McKenzie Meline, and Beth Harn. 2023. Student Engagement Within Adolescent Reading Comprehension Interventions: A Systematic Literature Review. *Journal of Education* 203, 2 (2023), 258–268. doi:10.1177/00220574211032327

[2] Rémi Cardon and Adrien Bibal. 2023. On Operations in Automatic Text Simplification. In *Proceedings of the Second Workshop on Text Simplification, Accessibility and Readability*. INCOMA Ltd., Varna, Bulgaria, 116–130. https://aclanthology.org/2023.tsar-1.12/

[3] Jiangjie Chen, Xintao Wang, Rui Xu, Siyu Yuan, Yikai Zhang, Wei Shi, Jian Xie, Shuang Li, Ruihan Yang, Tinghui Zhu, Aili Chen, Nianqi Li, Lida Chen, Caiyu Hu, Siye Wu, Scott Ren, Ziquan Fu, and Yanghua Xiao. 2024. From Persona to Personalization: A Survey on Role-Playing Language Agents. *Transactions on Machine Learning Research* (2024). https://openreview.net/forum?id=xrO70E8UIZ

[4] Yi-Pei Chen, Noriki Nishida, Hideki Nakayama, and Yuji Matsumoto. 2024. Recent Trends in Personalized Dialogue Generation: A Review of Datasets, Methodologies, and Evaluations. In *Proceedings of the 2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation (LREC-COLING 2024)*. ELRA and ICCL, Torino, Italia, 13650–13665. https://aclanthology.org/2024.lrec-main.1192/

[5] Zhengyu Hu, Jianxun Lian, Zheyuan Xiao, Max Xiong, Yuxuan Lei, Tianfu Wang, Kaize Ding, Ziang Xiao, Nicholas Jing Yuan, and Xing Xie. 2025. Population-Aligned Persona Generation for LLM-based Social Simulation. arXiv:2509.10127 [cs.CL]

[6] Ludivine Javourey-Drevet, Stephane Dufau, Thomas François, Núria Gala, Jacques Ginestié, and Johannes Ziegler. 2022. Simplification of literary and scientific texts to improve reading fluency and comprehension in beginning readers of French. *Applied Psycholinguistics* 43 (01 2022), 1–28. doi:10.1017/S014271642100062X

[7] Ang Li, Haozhe Chen, Hongseok Namkoong, and Tianyi Peng. 2025. LLM Generated Persona is a Promise with a Catch. arXiv:2503.16527 [cs.CL]

[8] Xun Liang, Hanyu Wang, Yezhaohui Wang, Shichao Song, Jiawei Yang, Simin Niu, Jie Hu, Dan Liu, Shunyu Yao, Feiyu Xiong, and Zhiyu Li. 2024. Controllable Text Generation for Large Language Models: A Survey. arXiv:2408.12599 [cs.CL]

[9] Joy Jia Yin Lim, Daniel Zhang-Li, Jifan Yu, Xin Cong, Ye He, Zhiyuan Liu, Huiqin Liu, Lei Hou, Juanzi Li, and Bin Xu. 2025. Learning in Context: Personalizing Educational Content with Large Language Models to Enhance Student Learning. arXiv:2509.15068 [cs.HC]

[10] Pedro Henrique Luz de Araujo, Paul Röttger, Dirk Hovy, and Benjamin Roth. 2025. Principled Personas: Defining and Measuring the Intended Effects of Persona Prompting on Task Performance. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 26845–26874. doi:10.18653/v1/2025.emnlp-main.1364

[11] Mounica Maddela, Fernando Alva-Manchego, and Wei Xu. 2021. Controllable Text Simplification with Explicit Paraphrasing. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. Association for Computational Linguistics, Online, 3536–3553. doi:10.18653/v1/2021.naacl-main.277

[12] Louis Martin, Éric de la Clergerie, Benoît Sagot, and Antoine Bordes. 2020. Controllable Sentence Simplification. In *12th Language Resources and Evaluation Conference*. European Language Resources Association, Marseille, France, 4689–4698. https://aclanthology.org/2020.lrec-1.577/

[13] OpenAI. 2026. GPT-5.4. https://openai.com/it-IT/index/introducing-gpt-5-4/ Accessed: March 2026.

[14] Emily Reif, Daphne Ippolito, Ann Yuan, Andy Coenen, Chris Callison-Burch, and Jason Wei. 2022. A Recipe for Arbitrary Text Style Transfer with Large Language Models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*. Association for Computational Linguistics, Dublin, Ireland, 837–848. doi:10.18653/v1/2022.acl-short.94

[15] Joni Salminen, Danial Amin, and Bernard J. Jansen. 2025. Using AI for User Representation: An Analysis of 83 Persona Prompts. In *IEEE/ACS 22nd International Conference on Computer Systems and Applications (AICCSA)*. IEEE, 1–8. doi:10.1109/aiccsa66935.2025.11315350

[16] William Schreiber, Jules White, and Douglas C. Schmidt. 2026. Toward A Pattern Language for Persona-Based Interactions with LLMs. In *31st Conference on Pattern Languages of Programs, People, and Practices* (Columbia River Gorge, WA, USA) *(PLoP '24)*. The Hillside Group, USA, 19 pages. doi:10.64346/PLoP2024p27

[17] Chenkai Sun, Ke Yang, Revanth Gangi Reddy, Yi Fung, Hou Pong Chan, Kevin Small, ChengXiang Zhai, and Heng Ji. 2025. Persona-DB: Efficient Large Language Model Personalization for Response Prediction with Collaborative Data Refinement. In *Proceedings of the 31st International Conference on Computational Linguistics*. Association for Computational Linguistics, Abu Dhabi, UAE, 281–296. https://aclanthology.org/2025.coling-main.20/

[18] Kimberly Truong, Riccardo Fogliato, Hoda Heidari, and Steven Wu. 2025. Persona-Augmented Benchmarking: Evaluating LLMs Across Diverse Writing Styles. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, Suzhou, China, 22676–22709. doi:10.18653/v1/2025.emnlp-main.1155

[19] Yu-Min Tseng, Yu-Chao Huang, Teng-Yun Hsiao, Wei-Lin Chen, Chao-Wei Huang, Yu Meng, and Yun-Nung Chen. 2024. Two Tales of Persona in LLMs: A Survey of Role-Playing and Personalization. In *Findings of the Association for Computational Linguistics: EMNLP 2024*. Association for Computational Linguistics, Miami, Florida, USA, 16612–16631. doi:10.18653/v1/2024.findings-emnlp.969

[20] Tatsuya Zetsu, Tomoyuki Kajiwara, and Yuki Arase. 2022. Lexically Constrained Decoding with Edit Operation Prediction for Controllable Text Simplification. In *Proceedings of the Workshop on Text Simplification, Accessibility, and Readability (TSAR-2022)*. Association for Computational Linguistics, Abu Dhabi, United Arab Emirates (Virtual), 147–153. doi:10.18653/v1/2022.tsar-1.13

[21] Mingqian Zheng, Jiaxin Pei, Lajanugen Logeswaran, Moontae Lee, and David Jurgens. 2024. When "A Helpful Assistant" Is Not Really Helpful: Personas in System Prompts Do Not Improve Performances of Large Language Models. In *Findings of the Association for Computational Linguistics: EMNLP 2024*. Association for Computational Linguistics, Miami, Florida, USA, 15126–15154. doi:10.18653/v1/2024.findings-emnlp.888