Generative AI as Mediator of Interpersonal Communication: Effects on Social Dynamics, Linguistic Norms, and Relational Authenticity

Vanja Stojković*

Deputy Director, National Employment Service of Republic Serbia

Abstract: The integration of generative artificial intelligence into everyday interpersonal communication - through AI-assisted writing, real-time translation, automated summarization, and conversational suggestion tools - represents a structural transformation in the mediation of human-human communication. Unlike earlier communication technologies that mainly transmitted or stored human-generated content, generative AI actively co-produces communicative output through a system architecture that includes user interfaces, prompt and context management, large language model inference, personalization modules, disclosure controls, and post-generation human editing. This paper examines AI-mediated interpersonal communication (AIMIC) as both a social phenomenon and an implementable communication system. It analyzes three interconnected effects: linguistic convergence, redistribution of communicative labor, and challenges to relational authenticity. The revised AIMIC framework links these effects to concrete implementation layers in email assistants, messaging platforms, translation tools, collaborative writing systems, and organizational communication software. The paper argues that responsible AIMIC design requires preserving human communicative agency, making AI involvement contextually transparent, monitoring linguistic and relational impacts, and embedding governance mechanisms into the communication pipeline. The paper concludes with practical examples, evaluation recommendations, and design principles for platform developers, educators, organizations, and policymakers.

Keywords: AI-mediated communication, Interpersonal communication, Generative AI, Linguistic norms, Relational authenticity, Communicative labor, Large language models, Social dynamics, CSCW, Digital communication, AI communication systems.

1. INTRODUCTION

Communication technologies have always shaped the social relationships they mediate. The telegraph compressed distance, the telephone changed expectations of presence and availability, and social media restructured the attention economy and the topology of interpersonal networks. Generative AI introduces a different kind of mediation because it does not merely transmit human-produced messages; it participates in producing the message itself.

This shift from channel mediation to content co-production changes the communicative status of everyday exchanges. A user may ask an AI assistant to draft an email of apology, rewrite a message to sound warmer, translate a sensitive statement into another language, summarize a long chat thread, or propose a reply in a professional messaging platform. In each case, the AI system becomes part of the communicative act, even when the final message is sent under the human user's name.

The manuscript therefore treats AI-mediated interpersonal communication (AIMIC) as an emerging socio-technical system rather than only as a theoretical topic in communication studies. The central contribution is the AIMIC framework, which links social outcomes - linguistic convergence, communicative labor redistribution, and authenticity challenges - to implementable system components such as prompt construction, language model generation, personalization, human-in-the-loop editing, disclosure mechanisms, and audit or provenance controls. This technical anchoring strengthens the paper's relevance for AI communication system design.

The paper is organized as follows. Section 2 outlines the research approach and scope. Section 3 reviews the theoretical background. Section 4 presents the AI communication system architecture underlying AIMIC and introduces the conceptual framework diagram. Section 5 analyzes the main social effects of AI-mediated interpersonal communication. Section 6 presents practical real-world examples. Section 7 discusses governance and evaluation recommendations, and Section 8 concludes.

2. RESEARCH APPROACH AND SCOPE

This paper is a conceptual and analytical review. It synthesizes literature from computer-mediated communication, sociolinguistics, human-computer interaction, AI ethics, and large language model research. Rather than presenting a new empirical dataset, it develops a structured framework for interpreting how generative AI changes interpersonal communication when integrated into operational communication systems.

*Address correspondence to this author at the Deputy Director, National Employment Service of Republic Serbia;
E-mail: vanjastojkovic988@gmail.com

Received on 20-04-2026 Accepted on 29-05-2026 Published on 05-07-2026

E-ISSN: 2979-1081/26

The analysis follows three steps. First, it identifies core communication-theoretical constructs affected by generative AI: communicative agency, relational transparency, linguistic impact, communicative labor, and authenticity. Second, it links these constructs to AI system implementation layers, including user interface design, prompt and context engineering, model output generation, personalization, translation, summarization, disclosure, and logging. Third, it derives governance and evaluation recommendations that can be used by platform designers, educators, organizations, and policymakers.

The scope is limited to AI systems that mediate human-human communication. The paper does not focus on purely human-AI companion relationships, fully autonomous customer-service bots, or general-purpose productivity tools unless they intervene in interpersonal message creation, translation, summarization, or response suggestion.

3. THEORETICAL BACKGROUND

3.1. Communication Mediation Theory

Medium theory argues that communication media are not neutral conduits but active shapers of social organization and cognitive style (Innis, 1951; McLuhan, 1964). Computer-mediated communication theory refined this insight by showing how asynchronicity, persistence, searchability, scalability, and platform affordances alter interpersonal relationship formation and maintenance (Walther, 1992; Herring, 2007).

Generative AI adds a new mediating layer. Earlier communication platforms changed the channel and temporal structure of communication, while generative AI changes the origination and formulation of content. It may select words, determine tone, restructure arguments, and translate meanings across languages. This makes AI-mediated interpersonal communication qualitatively different from conventional digital communication.

3.2. Communicative Labor and Relational Investment

Communicative labor refers to the cognitive, emotional, and temporal effort needed to produce a communicative act (Baym, 2015; Crawford, 2021). This labor is socially meaningful. A carefully written message often signals care, attention, and relational investment, while a generic or template-like message may signal distance or low commitment.

AI writing assistance disrupts this signaling mechanism. A message that appears thoughtful and personalized may have required only minimal human effort if generated by a model. Conversely, AI assistance may help people with language barriers, writing difficulties, disability-related constraints, or time pressure communicate more effectively. The central design challenge is therefore not whether AI assistance should exist, but whether it augments human agency or displaces human communicative effort in ways that mislead recipients.

3.3. Authenticity in Mediated Communication

Authenticity in communication concerns the perceived congruence between the expressed message and the sender's genuine intention, identity, and relational stance (Goffman, 1959; Hecht, 1993). All mediated communication involves some performance and adaptation, but AI co-authorship creates a new ambiguity: the message may be socially received as the sender's voice even when substantial linguistic and emotional formulation was produced by a system.

This ambiguity is especially important in emotionally sensitive settings such as apologies, condolences, praise, professional recommendations, conflict resolution, and intimate communication. In such contexts, the relational meaning of the message often depends not only on semantic content but also on the effort and vulnerability perceived behind it.

4. AIMIC AS AN AI COMMUNICATION SYSTEM ARCHITECTURE

To strengthen the technical perspective, AIMIC can be understood as a layered AI communication system. At the front end, the user interface captures a communicative intention: drafting an email, rewriting a message, translating content, summarizing a conversation, or generating a suggested reply. A context-management layer then packages relevant information, such as recipient identity, prior thread content, relationship type, language preferences, institutional templates, and user-selected tone. The AI mediation layer applies one or more models, typically an LLM, machine translation model, summarization model, or tone classifier. A governance layer then checks the output for policy compliance, disclosure requirements, manipulation risk, and provenance logging before the user edits and sends the message.

This architecture matters because the social effects of AIMIC are not created by the model alone. They emerge from the full pipeline: what options the interface makes easy, what context is supplied to the model, whether the system encourages human editing, whether AI involvement is disclosed, how outputs are logged, and whether recipients can recognize or contest AI-mediated content. The AIMIC framework therefore links social theory to implementation design.

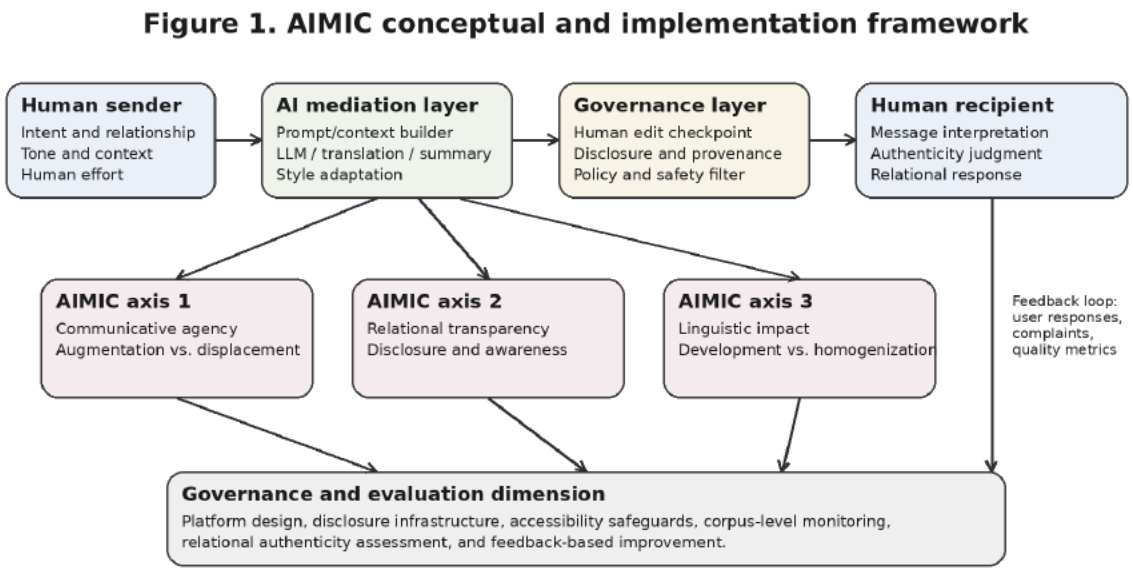

Figure 1: Conceptual and implementation model of the AIMIC framework.

4.1. Implementation Layers and Design Responsibilities

The AIMIC pipeline can be implemented differently across communication platforms, but the main layers are common. Table **1** maps the system components to design responsibilities and communication risks. This mapping is intended to help designers translate the conceptual framework into concrete interface and governance requirements.

4.2. Implementation Models

Three implementation models are especially relevant. The first is assistive composition, where AI proposes drafts or rewrites but the user remains the final author. The second is mediated translation and summarization, where AI transforms the language or length of a human-generated message while preserving its intended meaning. The third is semi-automated response generation, where AI recommends replies based on prior conversation context, organizational rules, or platform-specific tone settings.

Each model has a different risk profile. Assistive composition raises authorship and effort-signaling concerns; translation raises semantic fidelity and cultural nuance concerns; summarization raises omission, framing, and accountability concerns; automated reply suggestion raises relational authenticity concerns because recipients may attribute

Table 1: Technical Mapping of the AIMIC Framework to AI Communication System Components

System Layer	Typical Implementation Component	AIMIC Responsibility	Main Risk if Unmanaged
User interface	Email composer, chat window, translation panel, reply suggestion tool	Make AI assistance visible and keep the user actively involved in final authorship.	Uncritical sending of AI-generated content.
Context builder	Prompt template, recipient metadata, thread history, tone selector	Use context only when relevant and avoid hidden manipulation of relational tone.	Over-personalized or contextually misleading outputs.
AI mediation engine	LLM, translation model, summarizer, style rewriter	Support expression, clarity, accessibility, and cross-language communication.	Stylistic homogenization and displacement of human voice.
Human edit checkpoint	Review screen, tracked edits, comparison with original draft	Preserve communicative agency and encourage meaningful human revision.	Loss of authorship and reduced communicative skill development.
Disclosure and provenance	AI-use label, metadata, audit log, optional recipient notice	Enable context-sensitive transparency about AI involvement.	Authenticity misattribution and disclosure norm confusion.
Evaluation loop	User feedback, recipient response metrics, corpus-level style monitoring	Measure agency, transparency, linguistic diversity, and relational outcomes.	Optimization only for speed or engagement rather than communication quality.

emotional or professional attentiveness to the sender even when the response was largely generated by the platform.

5. EFFECTS OF AI-MEDIATED INTERPERSONAL COMMUNICATION

5.1. Linguistic Convergence and Stylistic Homogenization

One documented effect of large-scale AI writing assistance is convergence toward a shared AI-influenced style. Because LLM outputs reflect statistical regularities in training corpora, they often favor certain sentence structures, hedging patterns, transitional phrases, and politeness conventions (Ippolito *et al*., 2020; Guo *et al*., 2023). As AI-assisted communication scales, these patterns may enter professional and interpersonal communication across communities that previously exhibited more diverse stylistic profiles.

From a system-design perspective, convergence is not inevitable. Platforms can provide style-preserving rewriting, user-specific voice controls, comparison views between user draft and AI suggestion, and diversity-aware evaluation metrics. These technical choices determine whether the system expands expressive capacity or standardizes communication around a narrow model style.

5.2. Communicative Labor and Relational Investment Signaling

AI writing tools redistribute communicative labor. They can reduce the effort required for routine messages, lower barriers for non-native speakers, and support users with dyslexia or other communication difficulties. At the same time, they may weaken the connection between visible message quality and actual sender effort. This is most problematic when recipients interpret polished AI-assisted messages as evidence of personal attention or emotional investment.

The distinction between access-enabling use and investment-displacing use is central. A translation tool that helps a worker communicate accurately in a second language may preserve or enhance agency, while a fully automated condolence message may undermine relational authenticity. Governance should therefore avoid one-size-fits-all restrictions and instead evaluate the purpose, stakes, relationship type, and degree of human involvement.

5.3. Relational Authenticity and the Disclosure Dilemma

Recipients often evaluate messages not only by content quality but also by perceived authorship and effort. Studies of AI-assisted communication indicate that awareness of AI co-authorship can reduce perceived authenticity, especially in close relationships and emotionally significant exchanges (Hancock *et al*., 2020; Jakesch *et al*., 2023).

This produces a disclosure dilemma. If AI assistance is undisclosed, recipients may over-attribute effort and emotional presence to the sender. If AI assistance is disclosed in every context, communication may become burdened by excessive labels that do not reflect meaningful differences in agency. The AIMIC framework therefore recommends context-sensitive disclosure: stricter norms for intimate, educational, legal, health, and high-stakes professional communication; lighter norms for routine scheduling, accessibility support, or low-stakes grammatical correction.

5.4. Communicative Agency and Skill Development

Habitual AI assistance may affect the development of communicative competence. Writing skill develops through effortful practice: selecting words, revising drafts, adapting to audience, and learning from feedback (Bereiter & Scardamalia, 1987). If AI tools repeatedly replace these activities, users may become more dependent on model-generated expression and less confident in unaided writing.

However, AI can also support development when designed as a tutor rather than a substitute. Explanatory rewrites, side-by-side comparisons, revision rationales, and user-controlled tone alternatives can help users expand their expressive repertoire. The same technical system can therefore produce either developmental or atrophying AIMIC, depending on interaction design.

6. REAL-WORLD APPLICATION EXAMPLES

The practical relevance of AIMIC can be illustrated through common communication tools.

6.1. AI-Assisted Email

In professional email, AI systems may draft responses, soften tone, shorten long explanations, or generate polite follow-up messages. The benefit is efficiency and clarity, especially in high-volume administrative or business contexts. The risk is that recipients may interpret an AI-polished email as evidence of personal care or detailed attention that the sender did not actually invest. A responsible email implementation should therefore include a human edit checkpoint, optional AI-use disclosure, and a comparison view showing how the AI changed the original message.

6.2. Messaging and Suggested Replies

Messaging platforms increasingly provide suggested replies such as short acknowledgments, scheduling confirmations, or supportive phrases. These features reduce friction, but they also blur the line between spontaneous interpersonal responsiveness and automated relational maintenance. In low-stakes coordination, such suggestions may be harmless; in emotionally sensitive conversations, the system should avoid generating overly intimate or emotionally loaded replies without user reflection.

6.3. Translation Tools in Interpersonal Communication

Real-time translation tools enable communication across linguistic boundaries and can strengthen inclusion in multilingual workplaces, education, migration services, and public administration. Their AIMIC risk is not merely technical translation error but relational distortion: tone, politeness, humor, and cultural nuance may shift across languages. Translation systems should therefore provide tone indicators, ambiguity warnings, and user review options for emotionally or legally sensitive messages.

6.4. Conversation Summarization

AI summarization of meetings, chat threads, or long email chains can reduce cognitive load and improve organizational memory. At the interpersonal level, however, summarization may omit uncertainty, disagreement, humor, or emotional nuance that shaped the meaning of the exchange. For high-stakes decisions, AI summaries should preserve provenance, identify omitted material categories, and allow participants to correct the summary before it becomes an institutional record.

7. AIMIC FRAMEWORK: DESIGN AND GOVERNANCE RECOMMENDATIONS

The AIMIC framework is organized around three analytical axes and one governance dimension. These dimensions should be treated as design variables rather than purely theoretical concepts.

7.1. Communicative Agency Axis

The communicative agency axis describes whether AI augments or displaces human expression. Augmenting systems help users express their intended meaning more clearly, accessibly, or efficiently. Displacing systems generate content that users transmit without meaningful engagement. Interface design should keep users in the augmentation region by requiring review, encouraging edits, and providing explanations for suggested changes.

7.2. Relational Transparency Axis

The relational transparency axis describes whether communication partners understand the role AI played in message production. Transparency should be calibrated to context. A grammar correction may require little disclosure, while an AI-generated apology, recommendation letter, assessment, or emotionally sensitive message may require explicit disclosure or at least provenance information.

7.3. Linguistic Impact Axis

The linguistic impact axis describes whether AI assistance expands or narrows communicative capacity. Developmental AIMIC exposes users to improved expression while preserving practice and learning. Atrophying AIMIC reduces effort so completely that users lose opportunities to develop voice, style, and audience awareness. Platforms should evaluate not only output quality but also long-term user agency and linguistic diversity.

7.4. Governance Dimension

The governance dimension addresses responsibilities of platform designers, organizations, educators, and regulators. Key measures include disclosure infrastructure, configurable human-review requirements, accessibility safeguards, protection of user voice, corpus-level monitoring of stylistic homogenization, and institutional rules for AI use in assessment, hiring, public services, and emotionally sensitive communication.

8. Evaluation Recommendations

Evaluating AIMIC requires more than standard usability testing. A system can be fast and convenient while still weakening authenticity, user agency, or linguistic diversity. Recommended evaluation indicators include:

- Communicative agency: percentage of AI-generated text edited by the user, user-reported sense of authorship, and availability of revision explanations.
- Relational transparency: recipient awareness of AI involvement, appropriateness of disclosure mechanisms, and context sensitivity of labels.
- Linguistic diversity: stylometric variation across users before and after AI assistance, frequency of model-like phrasing, and preservation of user voice.
- Authenticity perception: recipient ratings of sincerity, effort, and relational appropriateness in AI-assisted versus unaided messages.

- Access and equity: benefits for non-native speakers, users with writing difficulties, and users with disabilities, balanced against risks of dependency.
- Governance performance: auditability, logging, complaint pathways, and ability to correct or contest AI-mediated communication records.

9. ETHICAL CONSIDERATIONS

The ethical concern is not that AI assistance is inherently inauthentic, but that it can obscure the social conditions under which a message was produced. If users benefit from AI assistance for accessibility, clarity, or translation, the system can promote communicative justice. If users rely on AI to simulate care, attention, expertise, or emotional availability that they did not provide, the same system can distort relationships.

Platform incentives deserve particular attention. Systems optimized only for speed, engagement, and frictionless response generation may encourage more automation than is socially desirable. Responsible AIMIC design should introduce proportionate friction in sensitive contexts, preserve human authorship, and help communities develop context-specific norms for AI-mediated communication.

Educational contexts require additional caution. Students should learn to use AI as a support for revision, translation, and feedback, but not as a replacement for the formative labor through which writing, reasoning, and interpersonal communication skills develop.

10. CONCLUSION

Generative AI is becoming a structural feature of interpersonal communication. Its significance lies not only in convenience or productivity but in the fact that it changes who, or what, participates in producing the words through which people maintain relationships, express identity, coordinate work, and signal care.

The revised AIMIC framework offers a technically grounded and socially oriented model for analyzing this transformation. By linking communicative agency, relational transparency, and linguistic impact to concrete AI communication system components, the framework provides guidance for platform developers, organizations, educators, and policymakers. Responsible AIMIC design should preserve human agency, support accessibility, make AI involvement transparent where it matters, protect linguistic diversity, and treat relational authenticity as a design and governance concern rather than as a peripheral cultural issue.

REFERENCES

Anderson, B., *et al*. (2023). AI writing assistants and communicative labor: An empirical analysis of usage patterns and stylistic effects in professional contexts. Journal of Computer-Mediated Communication, 28(4).

Baym, N. K. (2015). Personal Connections in the Digital Age (2nd ed.). Polity Press.

Bereiter, C., & Scardamalia, M. (1987). The Psychology of Written Composition. Lawrence Erlbaum Associates.

Crawford, K. (2021). Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence. Yale University Press. https://doi.org/10.12987/9780300252392

Goffman, E. (1959). The Presentation of Self in Everyday Life. Anchor Books.

Guo, B., *et al*. (2023). How close is ChatGPT to human experts? Professional writing quality and stylistic convergence analysis. arXiv preprint arXiv:2305.01852.

Hancock, J. T., Naaman, M., & Levy, K. (2020). AI-mediated communication: Definition, research agenda, and ethical considerations. Journal of Computer-Mediated Communication, 25(1), 89-100. https://doi.org/10.1093/jcmc/zmz022

Hecht, M. L. (1993). 2002 - A research odyssey: Toward the development of a communication theory of identity. Communication Monographs, 60(1), 76-82. https://doi.org/10.1080/03637759309376297

Herring, S. C. (2007). A faceted classification scheme for computer-mediated discourse. Language@Internet, 4(1).

Innis, H. A. (1951). The Bias of Communication. University of Toronto Press.

Ippolito, D., *et al*. (2020). Automatic detection of generated text is easiest when humans are fooled. Proceedings of ACL 2020. https://doi.org/10.18653/v1/2020.acl-main.164

Jakesch, M., Hancock, J. T., & Naaman, M. (2023). Human heuristics for AI-generated language are flawed. Proceedings of the National Academy of Sciences, 120(11). https://doi.org/10.1073/pnas.2208839120

Kasirzadeh, A., & Gabriel, I. (2023). In conversation with artificial intelligence: Aligning language models with human values. Philosophy & Technology, 36(2). https://doi.org/10.1007/s13347-023-00606-x

McLuhan, M. (1964). Understanding Media: The Extensions of Man. McGraw-Hill.

Walther, J. B. (1992). Interpersonal effects in computer-mediated interaction: A relational perspective. Communication Research, 19(1), 52-90. https://doi.org/10.1177/009365092019001003

Weidinger, L., *et al*. (2022). Taxonomy of risks posed by language models. Proceedings of FAccT 2022. https://doi.org/10.1145/3531146.3533088

https://doi.org/10.31875/2979-1081.2026.02.09