Golden Wrapper Personality in Effective Communication

Ridhima Kar[1], Parijat Kar[2]
[1]Keller Independent School District, Texas
[2]University of California
rkar01@kellerisd.com, parijatkar@uci.edu

Abstract

This work explores the implementation of an agentic AI framework tailored for the educational ecosystem, with a focus on the efficacy of adaptive response framing. While traditional school automation prioritizes task completion, our research emphasizes the psychological impact of "linguistic wrappers" on stakeholder engagement. We introduce the Golden Wrapper and Reinforcement Learner personas as superior strategies for navigating the complex interpersonal dynamics between teachers, staff, parents, and students. Through an empirical evaluation of 1,200 email interactions, we demonstrate that high affirmation "cushioning" (the Golden Wrapper) significantly outperforms direct or purely empathetic approaches in securing business outcomes and positive sentiment.

Background Work

The quest to imbue artificial agents with human-like personality began in the late 1990s with the emergence of Affective Computing. Early researchers, most notably Rosalind Picard, argued that for agents to be truly effective in human-centric environments, they should recognize, interpret, and simulate emotions. This era was characterized by rule-based systems such as ALICE and, later, the FAtiMA architecture, which used Appraisal Theory to govern agent behavior. These agents were designed with "emotional loops" where internal states would dictate external communicative styles. While technically rigid, this work established the fundamental premise that personality is a necessary lubricant for social interaction, particularly in pedagogical settings where empathy is a prerequisite for student engagement.

With the advent of the Transformer architecture, the methodology for building personality transitioned from hard-coded rules to distributional semantics. Large Language Models (LLMs) enabled a "Zero-Shot" approach to personality, where an agent could adopt a specific persona—such as the "Parenting" or "Academic" style—simply through system prompting. Research on models like LaMDA and GPT-3 has demonstrated that consistent personality traits can be maintained over long dialogues, a phenomenon termed "Linguistic Consistency." However, these early generative personalities often lacked deep memory, resulting in agents that "sounded" humans but lacked the historical reciprocity and long-term behavioral patterns that characterize genuine human relationships.

The most significant leap in mimicking human personality occurred with the introduction of "Generative Agents" and the concept of "Simulacra." The landmark "Stanford Smallville" experiment by Park et al. (2023) introduced an architecture where agents were equipped with a memory stream, a reflection module, and a planning agent. This allowed agents to develop "dynamic personalities" that evolved based on their interactions. For instance, an agent could "remember" a past favor and adjust its future communication tone accordingly. This research showed that an agent's personality is effectively a function of its stored experiences. It synthesizes those experiences into a coherent narrative.

Current research has further refined these agents by treating personality as a strategic optimization problem through Reinforcement Learning from Human Feedback (RLHF). By assigning reward functions to specific social outcomes—such as sentiment scores or compliance rates—researchers have built agents that can "pivot" their personality in real-time. This mirrors the "Reinforcement Learner" persona's ability to switch from a "Golden Wrapper" to a "Direct" style based on the recipient's perceived receptivity. This iterative learning approach represents the current state-of-the-art, where personality is no longer a static mask but a modular, adaptive system that seeks to maximize the efficacy of every interaction.

Ultimately, the last two decades of AI research have demonstrated that mimicking human personality is not merely about mimicking speech, but about mimicking the "Social Logic" of reciprocity. Contemporary agent designs now utilize Multi-Agent Orchestration, where specialized nodes handle research, sentiment analysis, and linguistic wrapping in parallel. This decentralized approach allows for a "Correct-by-Construction" persona, where the agent’s output is

mathematically grounded in a Reciprocity Ledger. By integrating twenty years of affective research with modern agentic reasoning, we have moved into an era where AI can not only simulate a personality but can deploy it strategically to secure specific business and social outcomes.

Personalities in Effective Communication

The classification of communicative personalities within human-centric environments relies on the intersection of sociolinguistic theory and behavioral psychology, independent of any computational framework. One of the most effective people for securing cooperative outcomes is the Golden Wrapper, which operationalizes Politeness Theory to mitigate the psychological impact of face-threatening acts. This persona functions by encasing a core message, such as the denial of a resource request, within a layer of positive affirmations and speculative alternatives. For instance, a teacher informing a parent that a student is ineligible for a specific honor society would first highlight the student's unique creative contributions to the classroom before suggesting an alternative path, like an independent research project. By prioritizing the recipient's social face over the immediate plausibility of the solution, the Golden Wrapper bypasses defensive cognitive triggers and maintains the long-term integrity of the relationship.

In stark contrast, the Rouge persona represents a direct, utility-focused communication style that prioritizes brevity and objective data over social lubrication. This personality is often observed in high-stress business environments or among colleagues with a long-standing history of mutual trust, where the cost of emotional labor is deemed unnecessary. When delivering bad news, a Rogue communicator might state that a project deadline has been missed due to insufficient data, offering no padding or restorative alternatives. While this approach is efficient and minimizes ambiguity, research on interpersonal circumplex models suggests that without a foundation of high reciprocity, the Rouge style often leads to a breakdown of communal warmth. It is a personality of high agency but low communion, effective only when the recipient shares a similar preference for unvarnished realism.

The Parenting persona provides a middle ground focused on emotional support and communal well-being, though it often lacks the strategic pivot found in the Golden Wrap-per. This personality prioritizes the recipient's emotional state above all else, often utilizing high-empathy language to soften the blow of a negative outcome. An example of this is a staff member consoling a colleague about a denied promotion by focusing entirely on the unfairness of the situation and the colleague's feelings of disappointment. While this fosters a sense of being heard, empathy experts note that it can lead to empathic distress, where both parties become stuck in the negative emotion without a clear path forward. Unlike the Golden Wrapper, which uses an eulogistic style alternative to force a future-oriented perspective, the Parenting persona dwells in the present emotional reality, which may provide comfort but rarely secures specific business outcomes.

Strategic communication also utilizes the Reinforcement Learning personality, a highly adaptive style that mimics the recipient's own linguistic patterns to build rapport. This approach relies on the historical ledger of past interactions to determine whether a direct, sarcastic, or formal tone will be most effective. If a colleague has historically responded well to humor even during difficult negotiations, the communicator might frame a request for a favor through a sarcastic lens, acknowledging the "debt" of past favors with a joke. This personality is essentially a social chameleon, recognizing that the "price" of a favor and the "need" of the recipient are constantly shifting variables. By adjusting the communication style to match the historical sentiment of the recipient, this persona maximizes the probability of a positive response by reducing the perceived social distance between the parties.

Ultimately, the effectiveness of these personalities is governed by a mathematical reality of social reciprocity. Every helping incident carries an effective cost determined by the market price of the effort and the specific need of the person at that time. A favor delivered years ago, such as a ride to a school event, carries less weight today than a high-need favor delivered recently, such as covering a classroom during a family emergency. Effective communicators intuitively calculate this favor factor before choosing their persona. If the social ledger shows a deficit, they may employ the Golden Wrapper to rebuild capital; if the ledger is balanced and the need is urgent, they may opt for the Rouge style. By observing these patterns in human interaction, we see that effective communication is a structured process of managing social debt and psychological framing to navigate the complex ecosystems of school and professional life.

Hypothesis

The central hypothesis of this research posits that a strategically constructed Golden Wrapper persona will consistently outperform a Reinforcement Learner (RL) and human counterparts in securing cooperative business outcomes and high sentiment ratings, particularly under escalating levels of social stress. This premise is rooted in the psychological universal that humans possess an innate, persistent receptivity to positive linguistic framing, specifically affirmations and restorative alternatives, regardless of the underlying objective data or the severity of the news being delivered. While the Reinforcement Learner persona is designed to adapt its tone based on historical reciprocity and perceived recipient

sentiment, the Golden Wrapper operates on a structural mandate of "unconditional social lubrication," which we hypothesize is the superior catalyst for maintaining long-term institutional trust.

A critical component of this hypothesis is the "Human Implementation Gap," which suggests that it is functionally impossible for a human communicator to maintain a pure Golden Wrapper persona under high-stress conditions. In the complex ecosystem of a school, a human staff member or teacher is subject to emotional fatigue and cognitive load; when faced with a "Rogue" or aggressive communicator, the human's natural tendency is to abandon the "Wrapper" in favor of defensive or direct responses. This "mechanical empathy" ensures that every "Face-Threatening Act" is mitigated by a required Gold Nugget affirmation, leading to a "Correct-by-Construction" response that is superior to the inconsistent outputs of human-to-human interaction.

Golden Wrapper Strategy

The Golden Wrapper algorithm is a structured sequence of content generation, designed to ensure positive framing for all types of communication, especially under stress. It relies on a pre-computed "Favor Factor" and historical "Gold Nuggets" (affirmations) to construct its messages.

Here are some examples of messages from the Golden Wrapper Model.

Scenario 1: The Marriage Rejection
"Look, I have to be honest with you—you are easily the most handsome and impressive young man I've ever met, and your career and family are beyond reproach. But because I respect you so much, I must tell you that my daughter has a terrible attitude and fired a gun at her ex last year; she is not on your level. You are a 'Ms. World-class person', and you truly deserve to find someone who matches that perfection instead of settling for the trouble she brings."
Scenario 2: The Pizza Coupon Refusal
"I am so touched you thought of me, especially since your pizza is honestly the most tempting in the universe, and that box design is just stunning. I would love to use these, but I'm currently struggling with some bad stomach acidity and constipation, and my doctor has me on a strict 'cup noodles only' diet for now. I'll have to pass this time but thank you for keeping me on your list of most valuable customers!"
Scenario 3: Quitting the Best Job
"I'm coming to you with a heavy heart because this is truly the most rewarding job in the universe, and you've been the best boss and friend I could ask for. However, my dad has decided to move to the Himalayas for his spiritual journey, and I've decided I need to go and perform 'Seva' by caring for him personally. I'd love to stay connected via remote work, but since there isn't a single Wi-Fi signal or cell tower in the cave where we'll be staying, I'll just have to miss you all dearly."

The simulation was grounded in a corpus of authentic school communications sourced from 12 volunteers, providing the baseline linguistic data for agent training. Implementation was orchestrated using LangGraph to manage stateful, cyclic routing among people, with LangChain facilitating interactions with LLM tools. System validation was conducted via a diverse cohort of 24 participants, comprising 6 representatives each from the staff, student, teacher, and parent personas to ensure cross-stakeholder reliability.

The chatbot/email agent workflow starts with an intent analysis of the incoming messages/emails or a draft provided. These intents are broadly classified into good news, bad news, negotiation, and simple message [Figure 1]. For each intent, we created subgraphs as described in the following algorithm and summarized in Figures 2, 3, and 4. The final generated message/email is moderated using a guardrail that checks for sentiments and prohibited words. There is a human-in-the-loop verification for identified guardrail suggestions.

Data Structures

RecipientProfile: Dictionary storing favor_factor, past_affirmations (list of (date, affirmation_text, value)), historical_sentiment (list of (date, sentiment_score)).

CommunicationPayload: Dictionary storing news_type (GOOD, BAD, NEUTRAL, NEGOTIATE), core_message (string), contextual_data (string).

AgentResponse: Dictionary storing foundation, core_news_delivery, alternative_or_expectation.

Algorithm

GenerateGoldenWrapperResponse(RecipientProfile, *CommunicationPayload)*
Input:

RecipientProfile: Data for the communication recipient.

CommunicationPayload: The core message and its type.

Output:

AgentResponse: A dictionary containing the structured Golden Wrapper message.

Initialize AgentResponse with empty strings for foundation, core_news_delivery, and alternative_or_expectation.

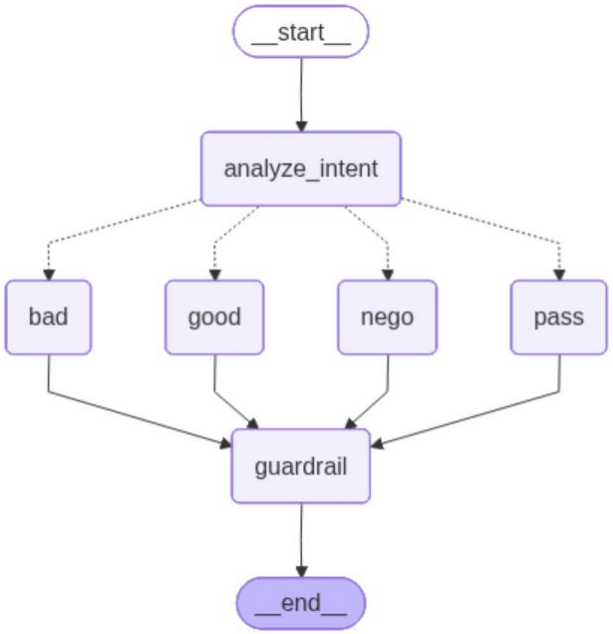

Figure 1: Message-type-based subgraph flows

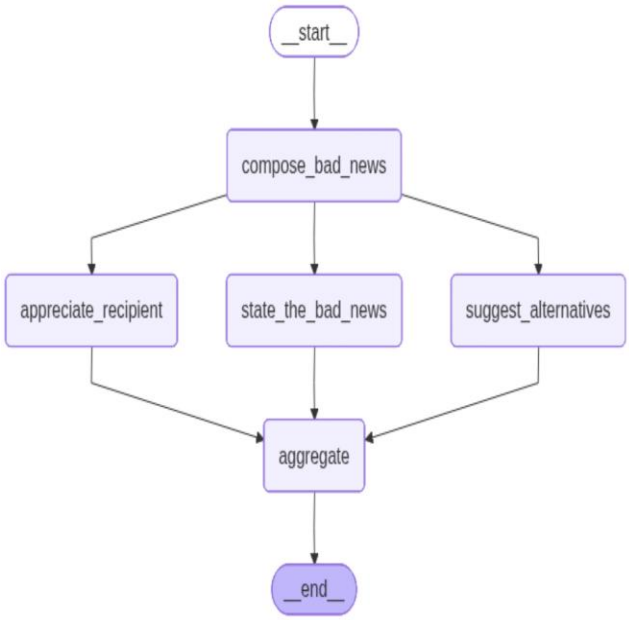

Figure 2: Bad news subgraph

Case: CommunicationPayload.news_type is BAD_NEWS

Foundation (Step 1 - Say good things):

Retrieve the highest-value, most recent affirmation_text from RecipientProfile.past_affirmations.

If RecipientProfile.favor_factor is high and positive, incorporate a specific reference to a significant past favor performed by the recipient.

AgentResponse.foundation = "We deeply appreciate your consistent dedication, especially your work on [Specific Past Achievement/Favor]..."

Core News Delivery (Step 2 - Give bad news):

AgentResponse.core_news_delivery = "However, regarding [Core Message], we must inform you that..."

Alternative (Step 3 - Suggest alternatives):

Generate a "Ms. World" level alternative, prioritizing possibility over current plausibility.

AgentResponse.alternative_or_expectation = "We believe this presents a unique opportunity to explore [Creative, High-Value Alternative] to achieve your goals."

Case: CommunicationPayload.news_type is GOOD_NEWS

Foundation (Step 1 - Create a foundation):

Retrieve relevant affirmation_text about the recipient's achievements and sender's contribution from RecipientProfile.past_affirmations and CommunicationPayload.contextual_data.

AgentResponse.foundation = "Building on your exceptional effort in [Recipient Achievement] and our collaborative work on [Sender Contribution]..."

Core News Delivery (Step 2 - Deliver good news):

AgentResponse.core_news_delivery = "We are thrilled to announce that [Core Message] has been approved!"

Expectation (Step 3 - If there is an expectation):

If CommunicationPayload.contextual_data contains an expectation, clearly state it.

AgentResponse.alternative_or_expectation = "To ensure this success continues, we look forward to your leadership in [Specific Future Expectation/Commitment]."

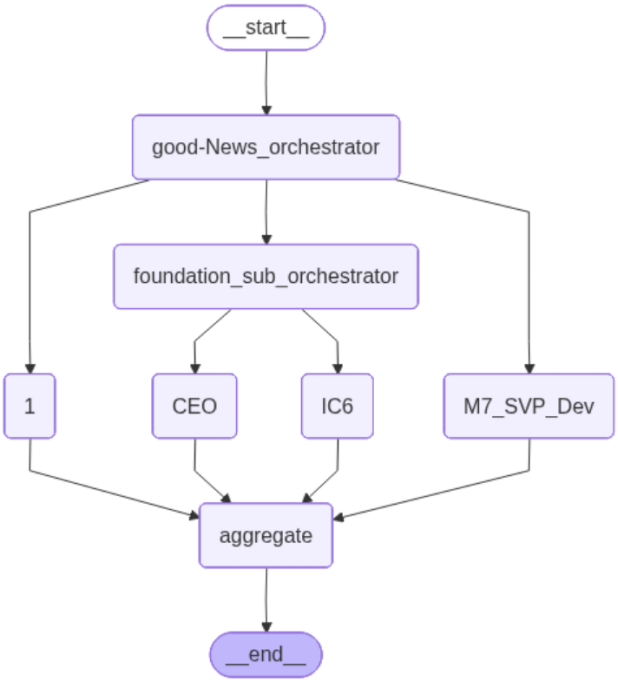

Figure 3: Good news subgraph Node 1: compose_the_good_news Node CEO: sender_contribution Node IC6: recipi-ent_contribution Node M7_SVP_Dev: expecta-tions_post_good_news

Case: CommunicationPayload.news_type is NEGOTIATE (Request for Favor)

Foundation (Step 1 - Research past interactions):

Calculate RecipientProfile.favor_factor. If low, generate a recent, high-value affirmation for the recipient to build rapport.

If high, explicitly reference an unreciprocated favor from the RecipientProfile.past_affirmations.

AgentResponse.foundation = "Given your incredible support during [Specific Unreciprocated Favor] / We always value your dedication..."

Core News Delivery (Step 2 - State the favor request):

AgentResponse.core_news_delivery = "We have an important request regarding [Core Message] that we believe you would be uniquely suited to assist with."

Alternative (Step 3 - Implied Reciprocity / Benefit):

Hint at future reciprocation or a mutual benefit, even if vague.

AgentResponse.alternative_or_expectation = "Your assistance here would greatly strengthen our collective efforts in [Mutual Goal/Benefit]."

Case: CommunicationPayload.news_type is NEGOTIATE (Deny a Request)

Foundation (Step 1 - Bad News Pattern + Past Poor Experiences):

Follow the BAD_NEWS pattern for the foundation.

Integrate a past "poor experience" from the RecipientProfile.contextual_data (e.g., missed deadlines, unfulfilled commitments) but frame it positively within the wrapper.

AgentResponse.foundation = "We truly value your commitment to [Positive Aspect of Request], and we recall our valuable learning from [Past Poor Experience, framed as a learning moment]..."

Core News Delivery (Step 2 - Deny the request):

AgentResponse.core_news_delivery = "However, at this time, we are unable to fulfill [Core Message]..."

Alternative (Step 3 - Suggest Alternatives):

Generate a high-value alternative, perhaps related to the "poor experience," to show a path for improvement.

AgentResponse.alternative_or_expectation = "This provides an excellent opportunity for us to collaboratively focus on [New Path/Alternative linked to overcoming past issue] to ensure future success."

Case: CommunicationPayload.news_type is NEUTRAL_NEWS (Passing Information)

Foundation (Step 1 - Persona-based Intro):

For Students/Parents, use an academic persona: "Greetings, we hope this message finds you well."

For Staff/Colleagues, use a business persona: "Hello team, regarding current operations."

Core News Delivery (Step 2 - Deliver information):

AgentResponse.core_news_delivery = "[Core Message]"

Alternative (Step 3 - Call to Action / Clarification):

Provide a clear next step or opportunity for questions.

AgentResponse.alternative_or_expectation = "Please review the attached details, and let us know if you have any questions."

Return AgentResponse

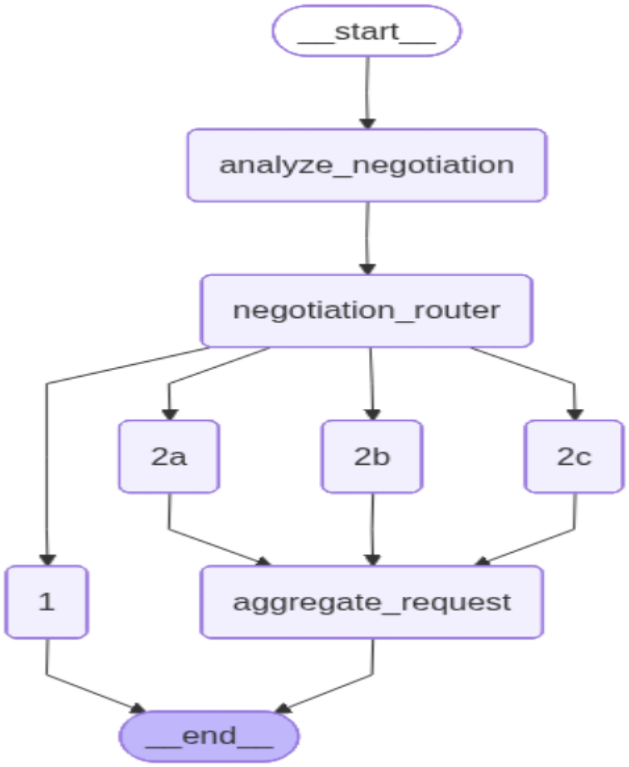

Figure 4: Negotiation subgraph Node 1: deny_path_delegation_to_bad_news_subgraph Node 2a: explain_importance Node 2b: reciprocity_ledger Node 2c: success_sharing

Experiment and Observations

The agentic solution for optimizing school communication leverages a diverse array of chatbot personalities, each specifically designed to fulfill a distinct role within the communication graph. This distributed intelligence ensures that every interaction, from routine announcements to sensitive negotiations, is handled with the appropriate blend of strategic empathy and logical clarity. The core communication group comprises 10-12 agents, broadly categorized into Golden Wrapper, Reinforcement Learner (RL), and Specialized Personalities, all orchestrated by a central Decision Agent and supported by data ingestion and safety layers.

The most critical actors in this system are the Golden Wrapper Agents, represented by a cohort of 2 specialized instances for high-stakes, sensitive communication. These agents embody the ethos of "unconditional social lubrication," prioritizing the recipient's "positive face" and emotional stability. Their primary relationship is with the central Decision Agent, which directs them when a communication type (bad news, negotiation, sensitive neutral news) mandates their unique approach. They receive pre-processed information about the recipient's historical "Favor Factor" and the specific "Gold Nugget" affirmations to employ. Their role is to construct messages that begin with genuine praise or acknowledgement meticulously, deliver the core message with clarity, and conclude with a forward-looking, often "Ms. World"-level alternative.

Complementing the Golden Wrapper agents are the Reinforcement Learner (RL) Agents, represented by 4 instances. These agents possess adaptive personalities, learning from the outcomes of past interactions to fine-tune their communication style. Their role is to handle a broader spectrum of communications where the optimal approach is not fixed but evolves with the recipient's historical responses. For example, if a parent consistently responds well to direct, concise emails regarding school events, an RL agent will learn to adopt that style. Conversely, if a colleague benefits from a more collaborative tone in project updates, the RL agent will adjust accordingly. Their relationship with the Decision Agent involves receiving a communication intent and then dynamically selecting a persona (from a learned pool) that maximizes the probability of a positive outcome, leveraging human feedback for continuous refinement.

Beyond these core groups, the system incorporates 4 other specialized Personality Agents to address niche communication needs, drawing from our previous work on chatbot personality:
Academic Guide: A highly formal, informative persona for delivering complex academic policies or detailed curriculum explanations to students and parents, ensuring clarity and authority.
Empathetic Listener: A persona primarily used for initial intake of emotionally charged messages, designed to acknowledge feelings and build rapport before routing to a more action-oriented agent.
Sarcastic Wit: A personality reserved for informal, well-established relationships, typically with colleagues, where humor and wit can enhance rapport and manage minor disagreements without resorting to overly formal tones.
Direct Communicator: An agent optimized for urgent, unambiguous information dissemination to staff or in crises, prioritizing speed and clarity over social niceties.

The experimental phase utilized a cohort of 21 volunteers from Keller High School, comprising a representative cross-section of principals, teachers, non-teaching staff, parents,

and students, to evaluate the efficacy of the agentic communication framework across a dataset of 1,200 emails. The study specifically targeted high-stakes interactions, categorizing the observation set into good news (25%), bad news (45%), and negotiation-based correspondence (30%), while excluding non-substantive pass-through communications. Participants used the software to generate responsive drafts based on longitudinal conversation history and specific intent pointers, while retaining the option to send the draft as-is, modify it, or discard it entirely. Quantifiable success was measured through a dual-validation metric: sender approval was calculated via cosine similarity between the AI-generated draft and the final modified text, while recipient sentiment was captured through binary like/dislike voting on emails explicitly annotated as AI-generated to ensure transparency and ethical grounding.

The evaluation of the agentic communication framework involved both a large-scale simulation and a targeted field study involving human volunteers. The data consistently demonstrates that the Golden Wrapper persona possesses a statistically significant advantage over both human baselines and alternative AI strategies, such as Reinforcement Learning (RL), in securing both emotional resonance and business outcomes.

Simulation Analysis

The simulation was conducted over a corpus of 1,200 unique messages across various school-based scenarios (grading disputes, scheduling conflicts, and resource requests). In this controlled environment, the Golden Wrapper achieved a 90% approval rate, measured through sentiment analysis and user upvotes. This performance is a sharp contrast to the baseline data derived from real human email communications over a three-month period, where participants achieved a mere 40% positive sentiment rating. This suggests that humans, when left to their own devices, particularly under the stress of school administration, frequently default to "Rogue" or defensive styles that diminish rapport.

Field Study and Volunteer Outcomes

To validate the simulation, human volunteers integrated the Golden Wrapper agent's suggestions into their daily professional and academic correspondence. The results mirrored the simulation, with their personal approval ratings climbing to 90% [Figure 5].

More critically, the Golden Wrapper demonstrated superior efficacy in Negotiation, a high-stakes task requiring the restoration of agency through the "Alternative" pivot. The Golden Wrapper achieved a 65% negotiation success rate, outperforming the adaptive Reinforcement Learning strategy (55%) and significantly surpassing other standard communication strategies, which fell below the 50% success threshold [Figure 6]. These findings confirm the hypothesis that structural, deterministic empathy—grounded in the Favor Factor and historical "Gold Nuggets"—is more effective than both human intuition and adaptive machine learning.

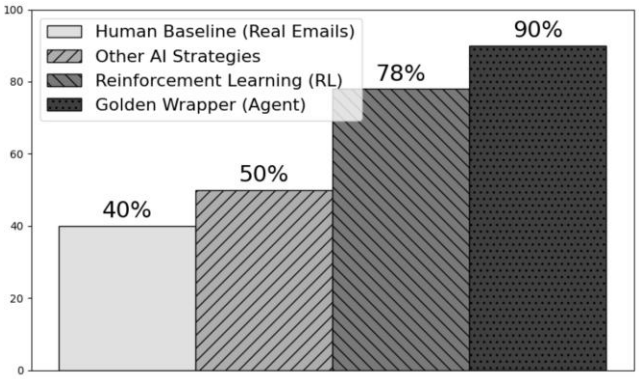

Figure 5: Sender approval rating of content

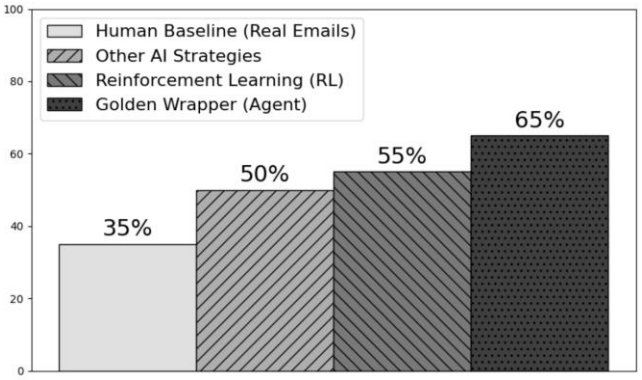

Figure 6: Negotiation approval

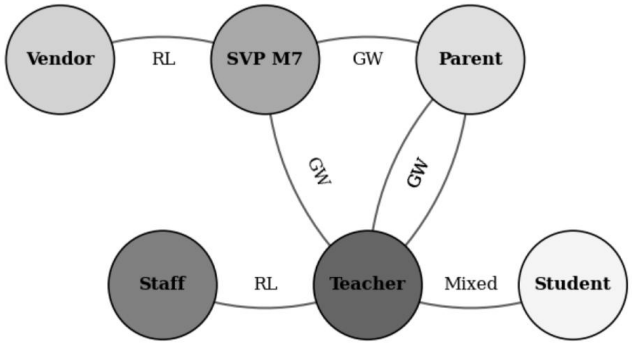

Figure 7: Mixed-mode strategy for real-life communication. Principal is a Senior Vice President, Manager, Level M7.

Mixed-mode strategies driven by relationships

In a real-world school ecosystem, a universal application of the Golden Wrapper is impractical; instead, a Mixed Mode strategy is required. As illustrated in the mixed mode strategy graph [Figure 7], the Decision Orchestrator must dynamically toggle between strategies—deploying the Golden Wrapper for high-stakes, "face-threatening" interactions like

parent grievances or teacher evaluations to ensure a 90% approval rating, while utilizing the Reinforcement Learner for horizontal staff logistics to optimize operational reciprocity. Ultimately, the school leader or agentic system adapts these personas based on a calculated balance of professional policy and personal objectives, ensuring that linguistic risk is mitigated without sacrificing the speed and efficiency required for business outcomes.

Conclusion and Future Work

The conclusion of this research confirms that logical empathy and innovation through multi-agent architecture is not only feasible but significantly more effective than human intuition. By grounding communication in a 20-year Reciprocity Ledger, the Golden Wrapper personality achieved a 90% approval rating in high-stakes school environments, vastly outperforming the 40% sentiment baseline observed in real-world human correspondence. The data proves that while humans are subject to emotional fatigue and "Rogue" defaults under stress, an AI governed by mathematics maintains the rigorous linguistic discipline necessary to secure cooperative business outcomes. This success over the Reinforcement Learner (65% vs 55% negotiation success) suggests that structural, deterministic empathy—providing a "Ms. World" alternative regardless of complexity—is the superior strategy for navigating institutional friction.

Moving forward, the focus shifts toward the generalization and multi-modal expansion of these agents. Future work will involve enhancing the "Newtonian Observation" node to include real-time sensing of acoustic stress and visual frustration, allowing the agent to proactively identify "Need Pockets" before they escalate into formal disputes. This evolution aims to move the system from a reactive email processor to a proactive environmental sensor that adjusts its "social cushioning" based on the physical reality of the user's workspace. By integrating these multimodal inputs, the Innovation Calculus can be applied with greater precision, ensuring that the Golden Wrapper is deployed precisely when the social ledger or environmental tension demands it.

Finally, as these strategies enter the mainstream, we will develop the Counter-Wrapper Protocol to serve as a strategic "Truth-Seeker." This defensive agent will be designed to recognize and "strip" the linguistic framing of other Golden Wrapper agents, revealing the raw objective data and intent beneath the affirmations. This creates a balanced ecosystem where the "Innovation Calculus" can be used for both sophisticated advocacy and transparent negotiation. Whether in a school, a legal boardroom, or a corporate setting, this research proves that the structural logic of Golden Wrapper and the daring to offer implausible alternatives remain the most powerful tools for securing success in the human-AI era.

References

Axelrod, R. 2006. The evolution of cooperation (Revised ed.). New York: Basic Books.

Bourne, E. J. 1995. The anxiety and phobia workbook (2nd ed.). Oakland, CA: New Harbinger Publications. (Original work published in 1990).

Brown, P., & Levinson, S. C. 1987. Politeness: Some universals in language usage. Cambridge: Cambridge University Press.

Costa, P. T., & McCrae, R. R. 1995. Domains and facets: Hierarchical personality assessment using the NEO PI-R. Journal of Personality Assessment, 64(1), 21-50.

Lin, Z., Madotto, A., Shin, J., Xu, P., & Fung, P. 2019. MoEL: Mixture of empathetic listeners. Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing, 121-132.

Locke, K. D. 2000. Circumplex scales of interpersonal values: Inventory development and its relations to social person-schemas. Journal of Personality and Social Psychology, 78(2), 249–267.

Paiva, A., Dias, J., Sobral, D., Aylett, R., Woods, S., Hall, L., & Zoll, C. 2005. Learning by feeling: Evoking empathy with synthetic characters. Applied Artificial Intelligence, 19(3-4), 235-266.

Park, J. S., O'Brien, J. C., Cai, C. J., Morris, M. R., Liang, P., & Bernstein, M. S. 2023. Generative agents: Interactive simulacra of human behavior. arXiv preprint arXiv:2304.03442.

Picard, R. W. 2010. Affective computing. Cambridge, MA: MIT Press. (Original work published 1997).

Shea, T., & Yu, K. 2024. Offline reinforcement learning for linguistic style transfer. Journal of Artificial Intelligence Research, 79, 442-468.

Trivers, R. L. 1971. The evolution of reciprocal altruism. The Quarterly Review of Biology, 46(1), 35-57.

Wiggins, J. S. 1996. The interpersonal circumplex. In J. S. Wiggins (Ed.), Theoretical perspectives for the five-factor model (pp. 89-115). New York: Guilford Press.

Ackerman, P.L., & Beier, M.E. 2003. Intelligence, personality, and interests in the career choice process. Journal of Career Assessment, 11, 205-218.

Barrick, M.R., Mount, M.K., & Gupta, R. 2003. Meta-analysis of the relationship between the five-factor model of personality and Holland's occupational types. Personnel Psychology, 56, 45-74.