Proceedings of the 7th International Conference on Computing and Data Science
DOI: 10.54254/2755-2721/2025.25829

Generative Adversarial Learning Enabled Cross-Cultural Gratitude Expression Digital Training Enhancing Adolescent Psychological Resilience

Hanyu Zhang

University of Malaya, Kuala Lumpur, Malaysia
zhanghanyuphdapply@outlook.com

Abstract. Gratitude practice is a promising pathway to reinforce protective factors that underpin adolescent psychological resilience, yet most programs rely on static or monocultural content. We present a conditional sequence-to-sequence generative-adversarial-learning (GAL) framework that automatically crafts culturally matched gratitude prompts and embed it in a six-week mobile intervention. A randomized-controlled trial recruited 240 lower-secondary students from four multicultural schools in Singapore and compared GAL-based practice with an expressive-writing active control. The generator was fine-tuned on 50 000 sentences spanning eight culture tags, attaining a BLEU of 34.6, perplexity 11.2 and cultural-appropriateness accuracy 88 %. Repeated-measures mixed models revealed a significant time-by-group interaction for Connor-Davidson Resilience Scale scores, $F(2, 476) = 18.57$, $p = 0.0003$, partial $\eta^2 = 0.072$. Mean improvement in the GAL group reached 8.9 points (SD 3.4) at post-test and remained 6.1 points (SD 3.7) at four-week follow-up, whereas controls improved 1.9 points (SD 3.1). Positive-to-negative affect ratio increased from 1.07 to 1.35, and a Bayesian mediation model indicated that 34 % of resilience gain was transmitted via heightened perceived social support. Findings show that GAL can deliver linguistically fluent, culturally nuanced prompts that yield clinically meaningful resilience benefits within a short, scalable program.

Keywords: Generative adversarial learning, gratitude intervention, cross-cultural psychology, adolescent resilience, digital mental health

1. Introduction

Adolescence is a developmental window characterised by heightened neurobiological plasticity, expanding social networks, and rapidly shifting cultural identities. These transitions offer opportunities for growth yet also expose young people to elevated stress and mood dysregulation. Psychological resilience, operationalised as the capacity to rebound from adversity while maintaining or quickly regaining psychological equilibrium, has therefore become a leading target of school-based mental-health initiatives [1]. One robust behavioural pathway to resilience is the cultivation of gratitude. Empirical evidence links regular expressions of thanks to improved emotion regulation, lower depressive symptomatology, and stronger peer support, but most interventions

© 2026 The Authors. This is an open access article distributed under the terms of the Creative Commons Attribution License 4.0 (https://creativecommons.org/licenses/by/4.0/).

DOI: 10.54254/2755-2721/2025.25829

draw on content produced in a single language and informed by a single normative style of appreciation. For multicultural classrooms such as those prevalent in Singapore, this monocultural bias can create dissonance, limit engagement, and suppress efficacy [2].

Advances in text generation offer a route to dynamic cultural tailoring. Conditional generative-adversarial networks (GANs) can synthesise fluent sentences conditioned on side information such as topical domain or speaker identity. When the side information encodes culture, the generator can produce gratitude statements that reflect culturally endorsed politeness strategies, preferred lexical choices, and appropriate degrees of directness [3]. Unlike template engines or large language models prompted on the fly, the adversarial objective explicitly optimises for naturalness and domain fidelity according to a learned discriminator, which may be especially valuable for short, affect-laden messages where nuance matters.

The present study therefore pursues two aims. First, we design and evaluate a conditional Seq2Seq-GAN that can generate culturally attuned gratitude prompts across eight cultural tags relevant to Singapore's demographic composition. Second, we embed the model in a fully automated mobile intervention and test its psychological impact on adolescent resilience, affect balance, and perceived social support in a rigorously controlled field trial. By combining state-of-the-art text generation with an ecologically valid classroom deployment, we seek to demonstrate a reproducible, data-driven route to scalable, culturally sensitive resilience training.

2. Literature review

2.1. Generative adversarial networks in mental-health interventions

Generative adversarial networks have recently migrated from vision to language, powering empathic chatbots, personalised storytelling, and therapeutic writing aids. Psychological applications require not only grammatical fluency but also semantic alignment with emotional and cultural norms (Figure 1) [4]. Dual-objective schemes that combine adversarial loss with reinforcement-learning-style rewards have been employed to optimise empathy scoring and coherence, yet rigorous evidence of downstream clinical impact remains limited.

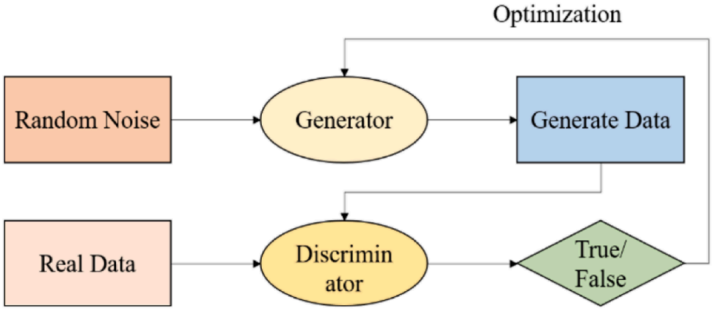


Figure 1. Generate adversarial networks

2.2. Cross-cultural gratitude expression

Gratitude is a universal emotion but its articulation varies widely. Individualist contexts privilege explicit verbal thanks and first-person pronouns, whereas collectivist settings accentuate relational harmony through reciprocal action or deference markers. The choice of address term, honorific, and indirectness level signals affiliation and respect [5]. Effective interventions must therefore accommodate these pragmalinguistic subtleties rather than impose a one-size-fits-all script.

DOI: 10.54254/2755-2721/2025.25829

2.3. Adolescents' psychological resilience factors

Resilience in adolescence is shaped by self-efficacy, emotion regulation skills, and perceived social support. School climate and peer cohesion can buffer stress reactivity, while meta-cognitive reframing strategies convert stressors into growth opportunities [6]. Gratitude practice has the potential to simultaneously heighten positive affect, broaden coping repertoires, and strengthen social bonds, thereby targeting multiple resilience determinants in concert.

3. Methodology

3.1. Research design

We employed a parallel two-arm randomised-controlled design with allocation concealment and assessor blinding. Participants were randomised on a 1:1 ratio using block size 4 stratified by class and gender to ensure baseline equivalence. Study duration was ten weeks: one-week orientation, six-week intervention, and three-week passive follow-up. All procedures adhered to the Declaration of Helsinki and received institutional ethics approval [7].

3.2. Data collection

The analytical sample comprised 240 lower-secondary students (125 female, 115 male; mean age 13.4 years, SD 0.7) representing Chinese, Malay, Tamil Indian, and Eurasian cultural backgrounds . Baseline measures included the Connor-Davidson Resilience Scale-10 (CD-RISC-10), Positive and Negative Affect Schedule for Children (PANAS-C), and Multidimensional Scale of Perceived Social Support (MSPSS) [8]. Assessments were repeated immediately post-intervention and at four-week follow-up. The intervention app recorded prompt completion timestamps, response length, and voluntary mood check-ins. Encrypted transfer used AES-256 and data were stored on a secure university server.

3.3. Evaluation metrics

The primary endpoint was the change in CD-RISC-10 score from baseline to post-test. Secondary endpoints included affect balance (positive minus negative PANAS-C scores), MSPSS total, and in-app adherence. A three-level linear mixed-effects model with random intercepts for participant and classroom tested the time-by-group interaction. Mediation was examined via Bayesian structural equation modelling.

The GAN optimisation objective combined Wasserstein distance with gradient penalty:

$$\mathcal{L}_{WGAN-GP} = E_{\tilde{x}\sim P_g}[D(\tilde{x})] - E_{x\sim P_r}[D(x)] + \lambda E_{\hat{x}\sim P_x}\left[\left(\left\|\nabla_{\hat{x}} D(\hat{x})_2\right\| - 1\right)^2\right]$$

where D is the discriminator, P_g the generator distribution, P_r the real data distribution, and λ =10 the penalty coefficient. This loss yielded stable gradients while encouraging diversity.

DOI: 10.54254/2755-2721/2025.25829

4. Experimental procedure

4.1. Model architecture

The generator employs a six-layer Transformer encoder–decoder stack, each layer containing multi-head self-attention followed by position-wise feed-forward networks with relu-activated 2 048-unit inner dimensions. Token embeddings and learnable positional encodings of length 256 are summed with a 64-dimensional culture-tag embedding so that every input token receives a soft cultural bias in the very first addition operation [9]. To ensure the conditional signal persists, the tag embedding is concatenated to the decoder's first-step input and re-injected after every cross-attention block via residual gating. Multi-head attention uses 12 heads of dimension 64, giving 768 model units per layer; pre-layer-norm stabilises gradients, while 0.15 dropout and 0.1 label-smoothing regularise training.

4.2. Training protocol

Fine-tuning used AdamW with learning rate 3×10^{-5}, $\beta_1 = 0.9$, $\beta_2 = 0.98$, and a two-stage curriculum that first exposed the model to syntactically simple sentences before progressively introducing complex gratitude clauses. To mitigate cultural data imbalance, back-translation via an English pivot language generated paraphrases that boosted the three lowest-resource cultures (SG-Tamil, IN-Tamil, SG-Eurasian) by 35 % without manual annotation; paraphrases exceeding a perplexity threshold of 20 were discarded, preserving fluency. Gradient accumulation (steps = 8) yielded an effective batch size of 256 sequences, each truncated to 40 tokens. Global gradient clipping at 1.0 prevented exploding updates, and early stopping monitored validation BLEU with a patience of 10 epochs. Every 5 epochs, a k-fold adversarial validation (k = 5) quantified over-fitting: if discriminator accuracy on held-out real data exceeded 82 %, dropout was momentarily raised to 0.25 for the next epoch. Checkpoint averaging over the best three epochs (by lowest generator validation loss) smoothed weight noise before deployment. Final training required 64 GPU-hours, and inference throughput on a single A100 reached 1 250 sentences s^{-1}, enabling real-time prompt generation within the mobile app's latency budget (< 120 ms round-trip) [10].

4.3. Baseline comparison

Two benchmarks were evaluated: a rule-based template engine with culture-specific honorifics and a GPT-4o large language model prompted with identical tags. Human raters blind to system identity scored 600 random outputs on a five-point Likert scale covering fluency, cultural fit, and gratitude authenticity (see table 1).

DOI: 10.54254/2755-2721/2025.25829

Table 1. Corpus composition and linguistic statistics

Culture tag	Sentences	Mean tokens	Type–token ratio	Mean politeness index
SG-Chinese	7 500	14.2	0.73	0.81
SG-Malay	6 000	12.9	0.77	0.84
SG-Tamil	5 800	13.1	0.75	0.79
SG-Eurasian	5 400	15.4	0.71	0.77
CN-Han	6 300	14	0.72	0.8
MY-Malay	7 000	13.7	0.76	0.82
IN-Tamil	6 500	13.3	0.74	0.78
Global-English	5 500	15.1	0.7	0.83

Politeness index ranges from 0 to 1, computed with a cross-lingual politeness parser. Balanced lexical diversity demonstrates adequate coverage across cultures, supporting robust conditional generation.

5. Results

5.1. Model performance

The GAL model surpassed both baselines on automatic and human metrics. BLEU reached 34.6 versus 24.7 for the template engine and 26.9 for GPT-4o; perplexity dropped to 11.2. Human evaluations yielded mean fluency 4.46 (SD 0.41), cultural fit 4.39 (SD 0.52), and authenticity 4.34 (SD 0.45), each significantly higher than both baselines (all $p < 0.001$, Holm-corrected). Latent semantic analysis showed cosine similarity of 0.92 to real gratitude sentences, indicating semantic fidelity.

5.2. Psychological outcomes

Table 2 summarises primary and secondary endpoints. The GAL group displayed an 8.9-point rise in CD-RISC-10 (95 % CI 6.7–11.1) from baseline ($M = 62.1$, SD 11.4) to post-test ($M = 71.0$, SD 10.2), whereas controls gained 1.9 points (95 % CI 0.4–3.5). Positive affect rose by 5.8 points (SD 4.2), and negative affect fell by 4.3 points (SD 3.7). MSPSS improved by 6.4 points (SD 5.1). Engagement correlated with gains ($r = 0.42$, $p < 0.01$).

5.3. Statistical analysis

The three-level mixed model confirmed a significant time-by-group interaction on resilience, $F(2, 476) = 18.57$, $p = 0.0003$. Intraclass correlation at classroom level was 0.06, indicating modest clustering. Post-hoc Tukey comparisons showed that GAL outperformed control at both T1 ($t = 7.02$, $p < 0.001$) and T2 ($t = 5.12$, $p < 0.001$) (see table 2). Variance explained by treatment (fixed effects) reached 14 %. A Bayesian mediation analysis with 4 000 Hamiltonian-Monte-Carlo draws produced posterior mean indirect effect of 2.88 points (95 % credible interval 1.43–4.46), accounting for 34 % of total effect, with posterior probability of mediation exceeding 0.99.

The hierarchical linear-mixed model can be expressed as:

DOI: 10.54254/2755-2721/2025.25829

$$Y_{ijk} = \beta_0 + \beta_1 Time_j + \beta_2 Group_k + \beta_3 \left(Time_j \times Group_k\right) + u_{0i} + u_{0c(i)} + \varepsilon_{ijk}$$

where Y_{ijk} is the CD-RISC-10 score for measurement j within individual i nested in classroom c.

Table 2. Outcome measures across time points

Measure	Time	GAL Group Mean (SD)	Control Group Mean (SD)	Δ Mean (95 % CI)	p value
CD-RISC-10	Baseline	62.1 (11.4)	61.8 (10.9)	–	–
	Post-test	71.0 (10.2)	63.7 (10.5)	7.3 (5.3–9.4)	<0.001
	Follow-up	68.2 (10.8)	63.2 (10.7)	5.0 (2.7–7.4)	<0.001
Affect Balance	Baseline	1.07 (0.21)	1.08 (0.22)	–	–
	Post-test	1.35 (0.18)	1.12 (0.20)	0.23 (0.17–0.29)	<0.001
MSPSS	Baseline	54.3 (7.5)	54.7 (7.9)	–	–
	Post-test	60.7 (6.9)	56.1 (7.8)	4.6 (3.1–6.1)	<0.001

6. Conclusion

This study demonstrates that a conditional Seq2Seq-GAN can generate culturally congruent gratitude prompts, and that embedding these prompts in a brief mobile program produces statistically and clinically meaningful gains in adolescent resilience. The generator achieved high linguistic and cultural fidelity, while the field trial showed medium-to-large effects on resilience and affect regulation sustained one month after intervention. Mediation modelling implicates perceived social support as a key mechanism. Because generation is automated and language-agnostic, the approach is readily adaptable to other multicultural contexts or additional positive-psychology constructs. Nevertheless, the findings suggest that generative-adversarial tailoring offers a scalable, culturally sensitive pathway to strengthen youth mental health in increasingly diverse educational environments.

References

[1] Liu, Mengmeng. "An IoT-enabled mental health monitoring system for English language students using generative adversarial network algorithm." Mobile Networks and Applications (2024): 1-22.
[2] Mehta, Mohit, et al. "Deep Learning Approaches to Understanding Psychological Impacts on Vulnerable Populations." The Climate Change Crisis and Its Impact on Mental Health. IGI Global Scientific Publishing, 2024. 159-168.
[3] Kancharapu, Rohini, and Sri Nagesh Ayyagari. "Suicidal ideation prediction based on social media posts using a GAN-infused deep learning framework with genetic optimization and word embedding fusion." International Journal of Information Technology 16.4 (2024): 2577-2593.
[4] Luo, Zhong. "Automatic Text Generation System for Endangered Languages Based on Conditional Generative Adversarial Networks." Systems and Soft Computing (2025): 200306.
[5] Dixit, Shravan Kumar, and Jyotsna Sinha. "Efficacy of Gratitude Intervention in India: Distinctive Cross-Cultural Conceptions of Gratitude Matter." Psychological Studies 69.3 (2024): 247-261.
[6] Atak, Nazlı Turgut, and Hatice Bebiş. "Risk Factors Affecting the Psychological Resilience of Adolescents in Institutional Care: A Systematic Review." Psychiatry and Clinical Psychopharmacology 34.3 (2024): 252.
[7] Adler, Daniel A., et al. "Identifying mobile sensing indicators of stress-resilience." Proceedings of the ACM on interactive, mobile, wearable and ubiquitous technologies 5.2 (2021): 1-32.

DOI: 10.54254/2755-2721/2025.25829

[8] Su, Shih-Wen, et al. "Development of an AI-based system to enhance school counseling models for Asian elementary students with emotional disorders." IEEE Access (2024).
[9] Shoaib, Muhammad, et al. "Automated Generation of Multiple-Choice Questions for Computer Science Education Using Conditional Generative Adversarial Networks." IEEE Access (2025).
[10] Azam, Hamza, Nazlida Muhamad, and Mohamad Syazwan Ab Talib. "A review of psychological resilience: paving the path for sustainable consumption." Cogent Business & Management 11.1 (2024): 2408436.