Proceedings of The 6th International Conference on Signal Processing and Machine Learning
DOI: 10.54254/2755-2721/2026.31158

# *Adaptive Recommendation Algorithms for Generative Art Therapy Content in Adolescent Trauma Interventions*

**Ziyi Zhu**

*Hangzhou Normal University, Hangzhou, China*
*rara481846778@gmail.com*

***Abstract.*** Adolescents exposed to family dysfunction, violence, and school bullying often show trauma-related symptoms and disengagement from class, and current digital interventions rely on static templates that do not handle trauma safety and individual differences well. This study designs an adaptive recommendation framework for generative art therapy that integrates content generation and tagging, a constrained contextual bandit policy, and behavioral and scale-based evaluation, and manages image, music, and narrative micro-tasks with structured labels of goal type, difficulty, and arousal. A school-based trial compares the adaptive policy with a static rule-based sequence. Results indicate higher task completion, deeper sessions, and higher next-day login in the adaptive group, together with larger improvements in emotion-regulation difficulty, self-efficacy, and trauma symptom scores and no increase in the proportion of "trigger/strong discomfort" sessions. The findings suggest that trauma-informed online learning policies can enhance engagement and psychological benefit in adolescent digital art therapy without an observable rise in safety incidents.

***Keywords:*** adolescent trauma, generative art therapy, adaptive recommendation, contextual bandits, trauma-informed safety constraints

## 1. Introduction

Adverse experiences such as early violence exposure, family dysfunction, and school bullying greatly increase the risk of trauma-related symptoms and learning difficulties in adolescents, which often appear as emotion dysregulation, attentional problems, and peer relationship breakdown across both mental-health and educational settings [1]. Art therapy, grounded in images, music, and narrative resources, can support emotional expression and self-integration under a safer symbolic distance and has been partially adopted in school-based support, but digital implementations are often limited to fixed task libraries or simple interest-based recommendation without fine-grained adaptation to momentary affect and learning pace or explicit computational safety constraints for generative content [2]. With the rapid advance of generative models, a central question for the intersection of educational psychology and computational intervention is how to integrate controllable content generation with adaptive recommendation so that task difficulty, expressive modality, and exposure intensity can be adjusted online under trauma-informed principles to enhance engagement and learning-related psychological outcomes [3]. In response, this study

DOI: 10.54254/2755-2721/2026.31158

proposes a generative art therapy content recommendation framework for adolescent trauma interventions and examines its system design, safety-constrained algorithm, and empirical evaluation to test its potential benefits and boundary conditions.

## 2. Literature review

### 2.1. Digital delivery mechanisms for adolescent trauma interventions

Digital delivery through mobile devices and online platforms extends the reach of services and embeds self-assessment, psychoeducation modules, and emotion-training exercises into everyday school and home contexts, but many programs still rely on one-size-fits-all content and fixed linear flows that do not account for differences in trauma history, executive functioning, or fluctuating motivation, which produces mismatched task load and early disengagement among high-risk or low-motivation users [4]. In educational technology, adaptive learning systems have demonstrated the potential to improve mastery and persistence through stratified difficulty, pacing adjustments, and personalized feedback, yet these mechanisms have not been systematically transferred into trauma-informed digital interventions and remain disconnected from algorithmic frameworks that jointly consider therapeutic benefit and safety, creating a critical gap between research on adaptation and practice in trauma care [5].

### 2.2. Psychological mechanisms of generative art therapy content

Art therapy theory suggests that media such as visual composition, color choice, musical rhythm, and narrative experimentation allow individuals to externalize emotions, reorganize fragments of traumatic memory, and test new self-positions within a low-threat perceptual environment, which helps restore a sense of control and coherence at a symbolic level and is particularly valuable for adolescents with limited verbal expression or strong defensive patterns [6]. Digital art therapy often relies on fixed templates and small asset libraries, which makes it difficult to sustain a fresh and developmentally attuned expressive space across repeated sessions, whereas generative models can produce a highly tunable spectrum of styles, themes, and complexity that can be aligned with goals such as graded exposure, resource activation, and strength-based work [7].

### 2.3. Evidence on online learning for adaptive recommendation in interventions

Research on recommender systems in education has shown that modeling learning trajectories and behavioral feedback with contextual bandits and reinforcement learning can optimize the sequence of resource delivery, reducing ineffective practice time and increasing subjective engagement while maintaining or improving overall performance, which provides an algorithmic basis for personalized ordering of intervention tasks [8]. In mental-health contexts, emerging work has started to use online learning to recommend modules such as meditation, behavioral activation, or sleep training based on self-reported mood and short-term responses, but many systems still treat clicks or completion as the main utility signal and rarely integrate symptom change, risk events, or ethical constraints into the decision process, and fairness across subgroups is often not examined [9]. For adolescent trauma interventions, algorithms must continuously balance the exploration of potentially beneficial content combinations against the avoidance of high-risk exposures, while also accounting for the distorting effects of school schedules, family support, and cultural norms on observed behavior, so online learning frameworks need reward designs, constraint strategies, and evaluation methods that are

DOI: 10.54254/2755-2721/2026.31158

interpretable from an educational-psychology perspective and acceptable from a clinical safety perspective.

## 3. Experimental methods

### 3.1. System architecture and the generation–tagging pipeline

The platform uses a front–back separation and microservice architecture. The web and mobile clients call backend REST APIs through an Nginx reverse proxy and an API gateway. The gateway writes user IDs, session IDs, and timestamps to a Kafka log queue and dispatches each request to the online recommendation service and the content service. The content service runs containerized diffusion and text generation models and receives a constraint vector of "goal type–difficulty–modality" through a gRPC interface to generate images, music, and narrative prompts. Generated outputs are stored in temporary object storage [10]. A separate tagging service calls sentiment models and visual arousal estimators, combines them with a keyword blacklist, assigns topic, goal, and arousal labels, writes tasks into a PostgreSQL task store, and caches high-frequency safe tasks in Redis. At each decision step, the recommendation service reads the last K task records and the current self-rated emotion from Redis, builds a fixed-length context vector, normalizes features, and feeds them into a contextual bandit policy that computes an upper bound of expected reward and a risk score for each arm. A safety filter removes arms whose risk exceeds the preset threshold and returns the ID of the selected task to the client. The client logs completion, dwell time, and post-session emotion scores via the gateway back to Kafka. An asynchronous learner updates policy parameters and the risk model from this log stream, which forms an end-to-end loop of log streaming, feature construction, policy decision, and feedback update.

### 3.2. Participants, study design, and measurement system

The study recruits 120 students aged 13–18 years from two regular high schools after standardized trauma screening and guardian informed consent. Stratified randomization based on the median trauma symptom score assigns 60 participants to an adaptive recommendation group and 60 to a static rule-based control group, with comparable distributions of gender, grade, and symptom levels.

The intervention lasts four weeks, with two online sessions per week and at least six micro-tasks per session. The first two sessions use only low-arousal resources to help users become familiar with the interface. From the third session on, the adaptive group uses the online learning policy, while the control group follows a fixed pre-defined sequence of goal type and difficulty. School psychologists monitor all sessions in real time and stop tasks in case of extreme reactions. Behavioral data include task completion rate, mean session depth, dwell time per task, and next-day voluntary login count. Psychological measures include the Chinese short form of the DERS emotion regulation difficulties scale, the GSES general self-efficacy scale, and a brief trauma symptom scale at baseline and post-intervention, plus 0–10 session-level distress ratings before and after each session, aligned with task sequences.

### 3.3. Constrained recommendation training and evaluation plan

The recommendation policy is modeled as a contextual multi-armed bandit. At task decision point ttt, a context vector $x_t$ and a candidate arm set $A_t$ are given [11]. The instantaneous reward as shown in Equation (1).

DOI: 10.54254/2755-2721/2026.31158

$$r_t = \alpha \tilde{e}_t + \beta \Delta \tilde{s}_t - \gamma \rho_t \quad (1)$$

Where $\tilde{e}_t$ is a normalized engagement score based on completion and dwell time, $\Delta \tilde{s}_t$ is the sign-reversed standardized change in distress within the session, and $\rho_t \in [0,1]$ is the risk score of the task in the current context.

Positive weights $\alpha,\beta,\gamma$ are selected by offline grid search. A linear UCB-style policy selects the arm as shown in Equation (2).

$$a_t = \underset{a \in A_t}{\operatorname{argmax}} \left( \hat{\mu}_a(x_t) + \lambda \sqrt{x_t^\top A_a^{-1} x_t} \right) \text{ s.t. } \hat{\rho}_a(x_t) \le \tau \quad (2)$$

Where $\hat{\mu}_a(x_t)$ is the estimated expected reward from a linear model for arm a, $A_a$ is the design matrix for that arm, $\lambda$ controls the width of the confidence bound, $\hat{\rho}_a x_t)$ is the output of a separate risk regression model, and $\tau$ is the risk threshold. After each task, the system updates $\hat{\mu}_a$ and $A_a$ using the observed reward and updates the risk model with samples labeled as "trigger" or "strong discomfort," which leads to joint learning of utility and safety within one framework.

## 4. Results

### 4.1. Evidence on recommendation utility and engagement gains

Among participants who complete all eight sessions, valid samples include 54 users in the adaptive group and 52 in the control group. Baseline completion rate and session depth show no significant difference between groups. After four weeks, mean task completion rate in the adaptive group is 78.4%, and in the control group is 65.1%; the difference is 13.3 percentage points, with $t(104)=3.12, p=0.002$. Mean session depth is 5.3 tasks in the adaptive group and 4.1 in the control group; the difference is 1.2 tasks, with $t(104)=2.87, p=0.005$, indicating longer task chains within similar session time. Next-day voluntary login rate is 36.0% in the adaptive group and 27.9% in the control group; the difference is 8.1 percentage points, with $\chi^2(1)=4.07, p=0.044$. Figure 1 shows a bar chart of completion rate and session depth by group, and the pattern is consistent across all three engagement indicators with limited overlap of confidence intervals.

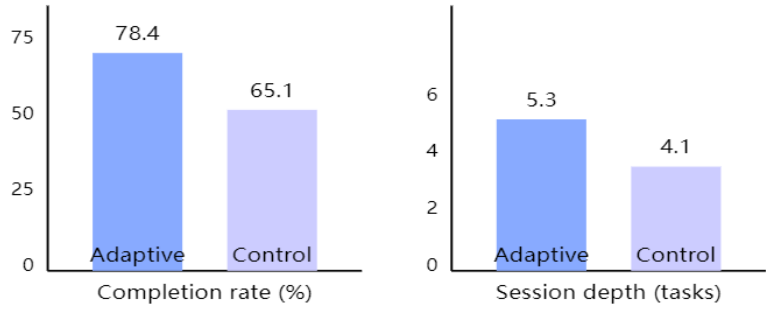


Figure 1. Completion rate and session depth by group

### 4.2. Psychological and educational outcomes with safety profile

For participants with complete pre–post scale data, valid samples include 51 users in the adaptive group and 49 in the control group. Lower DERS scores indicate fewer emotion-regulation

DOI: 10.54254/2755-2721/2026.31158

difficulties. Higher GSES scores indicate stronger self-efficacy. Lower trauma symptom scores indicate fewer symptoms. After four weeks, mean DERS score in the adaptive group decreases from 92.1 to 81.4, a change of −10.7 points. In the control group, mean DERS decreases from 91.7 to 86.9, a change of −4.8 points; the between-group difference in change is significant, with $t(98)=2.41, p=0.018$. Mean GSES in the adaptive group increases from 24.3 to 27.8, while in the control group it increases from 24.5 to 26.1; the difference in change yields $t(98)=2.07, p=0.041$. Mean trauma symptom score in the adaptive group decreases from 31.6 to 26.2, and in the control group from 31.9 to 29.4; the difference in change yields $t(98)=2.02, p=0.046$. The proportion of sessions labeled as "trigger" or "strong discomfort" is 1.3% for the adaptive group (3 sessions across eight rounds) and 2.7% for the control group (7 sessions across eight rounds), with Fisher's exact test $p=0.19$. Table 1 reports pre and post means and standard deviations for all three scales and the safety indicator.

Table 1. Pre–post comparison of main psychological scales and safety by group

| Scale | Group | Pre mean ± SD | Post mean ± SD | Change Δ |
|---|---|---|---|---|
| DERS total score | Adaptive | 92.1 ± 11.4 | 81.4 ± 12.0 | −10.7 |
| | Control | 91.7 ± 10.9 | 86.9 ± 11.5 | −4.8 |
| GSES total score | Adaptive | 24.3 ± 3.6 | 27.8 ± 3.9 | +3.5 |
| | Control | 24.5 ± 3.8 | 26.1 ± 3.7 | +1.6 |
| Trauma symptom total | Adaptive | 31.6 ± 7.2 | 26.2 ± 7.8 | −5.4 |
| | Control | 31.9 ± 7.5 | 29.4 ± 7.9 | −2.5 |
| Trigger/strong discomfort | Adaptive | — | 1.3% of sessions | — |
| | Control | — | 2.7% of sessions | — |

## 5. Discussion

Findings indicate that a constrained contextual bandit policy in a generative art therapy setting can increase behavioral engagement and practice intensity for adolescents within trauma-informed safety limits. Joint gains in completion rate, session depth, and next-day login suggest that fine-grained adjustment based on recent task history and momentary distress yields better alignment with each user's current load window and expressive preference than a static "goal by difficulty" schedule. Differences in changes on emotion-regulation difficulty, self-efficacy, and trauma symptom scales show that embedding both engagement and low-risk affective improvement into the reward function supports empirical approximation of tolerable challenge task combinations, rather than simply drifting toward the lowest arousal or purely entertaining content, while the proportion of "trigger/strong discomfort" sessions does not rise, which supports basic usability of the risk model and threshold constraints in a live log stream. Several limitations remain, including modest sample size, short follow-up, and possible novelty effects in the reward signal, and school context plus the presence of staff likely contribute additional buffering for safety and adherence, so future work should compare alternative algorithm variants and reward designs with stronger causal identification across longer periods and more diverse settings.

DOI: 10.54254/2755-2721/2026.31158

## 6. Conclusion

This study targets generative art therapy in adolescent trauma interventions and proposes an adaptive recommendation framework that combines log-based feature construction, a constrained contextual bandit policy, and a risk regression model for online selection and safety control of image, music, and narrative micro-tasks. In a four-week school pilot, the adaptive group shows larger gains than the static rule group in task completion, session depth, next-day logins, and changes in emotion-regulation difficulty, self-efficacy, and trauma symptom scores, with no increase in the proportion of trigger events. The framework provides a technical basis for larger multi-school trials of generative art therapy systems and an operational path for encoding trauma-informed principles into recommendation algorithms for educational psychology interventions.

## References

[1] Zhang, Bo, et al. "The effect of the art therapy interventions to alleviate depression symptoms among children and adolescents: a systematic review and meta-analysis." Clinics 80 (2025): 100683.
[2] Applewhite, Briana, et al. "Creative arts-based interventions for the improvement of PTSD symptoms in young people: a meta-analysis with a focus on non-Western populations." Nature Mental Health (2025): 1-17.
[3] Wang, Jiahua, et al. "Colors of the mind: a meta-analysis of creative arts therapy as an approach for post-traumatic stress disorder intervention." BMC psychology 13.1 (2025): 32.
[4] Blomdahl, Christina, and Angeliki Goulias. "Art therapy for adolescents with depression: feasibility and acceptability study in child and adolescent psychiatry." Art Therapy 42.1 (2025): 14-23.
[5] Wei, Shijuan, Angel Hor Yan Lai, and Howard Wing Hong Ho. "The Effectiveness of Art Therapy on Children and Adolescents with ASD: A Systematic Review of RCTs." Healthcare. Vol. 13. No. 22. 2025.
[6] Sheriff, Rebecca Syed, et al. "Online arts and culture for mental health in young people: a qualitative interview study." BMJ open 13.6 (2023): e071387.
[7] Sheriff, Rebecca J. Syed, et al. "Co-design of 'Ways of Being', a web-based experience to optimise online arts and culture for mental health in young people." BJPsych Bulletin 49.1 (2025): 41-48.
[8] Hong, Sun Hwa, et al. "Digital Mental Health Interventions for Adolescents: An Integrative Review Based on the Behavior Change Approach." Children 12.6 (2025): 770.
[9] Wani, Carolina, et al. "Digital mental health interventions for adolescents in low-and middle-income countries: scoping review." Journal of Medical Internet Research 26 (2024): e51376.
[10] Potts, Courtney, et al. "Digital Mental Health Interventions for Young People Aged 16-25 Years: Scoping Review." Journal of Medical Internet Research 27 (2025): e72892.
[11] Aguilera, Adrian, et al. "Effectiveness of a digital health intervention leveraging reinforcement learning: results From the Diabetes and Mental Health Adaptive Notification Tracking and Evaluation (DIAMANTE) Randomized Clinical Trial." Journal of medical Internet research 26 (2024): e60834.