

Redi-up: Human Feedback-Driven Preference Inference for Personalized Interior Redesign With AR Verification

Kim, Ian Tekhyun (School: Chadwick International School)

This research shows LLM-generated design can be customized by Deep Learning's Reinforcement learning with non-verbal(personal preference) human feedback. In space design, most of its applications ask users to describe their taste through keywords or short prompts. But aesthetic preference is inherently subjective — "modern" means something different to every person, and almost certainly different from what the AI assumes it to mean. The resulting design often mismatches user expectations, as it is shaped by the AI's interpretation of their words, not by their actual taste. Redi-up addresses this by learning user taste through human feedback rather than verbal description. After the user scans their room, the system presents sets of three redesigned interior images, from which the user selects one, generating pairwise preference signals from both selections and rejections. These signals are processed by a "style-classification ensemble model" trained on 20,000 interior images across 10 style classes. Then, the model updates a 10-dimensional preference vector that represents the user's aesthetic profile. This profile is then used to generate increasingly personalized image sets, forming an iterative refinement loop. The final design is converted into 3D and placed into the user's physical space using ARCore for real-world validation. The ensemble model achieved a test accuracy of 96.2%. The preference vector update magnitude decreased by 88.2% and stabilized by round 12.4, indicating that the system progressively converges on a user's definite aesthetic preference. Redi-up demonstrates that aesthetic taste can be precisely inferred from non-verbal human feedback alone — without prompts — closing the gap between AI-generated design and users' actual preferences

Awards Won:

Midjourney: Eight awards of \$2,500