Digital Intelligence PRESS

Journal of Discovery Core

https://doi.org/10.67541/jdc2610

*Research Article* **Volume 1, Issue 2 (Sep 2026)** *Open Access*

# Relation-Aware and Constraint-Guided Discrete Diffusion for Design Color Scheme Generation

**Mengyu Liu, Qiao He ***

*College of Art and Design, Beijing city university, Beijing, China*

***Corresponding author:** Qiao He. Email: qhe856006@gmail.com

## Abstract

The color collocation of the design scheme should meet multiple constraints such as functional semantics, visual harmony, accessibility and style consistency at the same time. The existing rule template, image extraction and condition generation methods have structural contradictions between constraint controllability and combination diversity, and the color relationship is always outside the generation process, only for post-output screening. In this paper, the problem is formulated as a sequential decision problem of discrete color combinations under multiple constraints, and the relation-aware diffusion generation framework ColorRel-Diff is proposed. The discrete color token sequence is used as the generation unit, and the mask discrete diffusion is used to make each denoising step explicitly masked discrete diffusion. The color relationship coding module is constructed, and the harmony and contrast relationships are encoded as the edge features of the relationship graph and aggregated into the relationship perception embedding through the graph attention. A constraining-guided denoising mechanism is designed to guide the sampling trajectory by the gradient of relational consistency score, and the feasibility mask is used to impose inviolable token-level restrictions on the hard constraints, and the grading softening strategy is used to deal with constraint conflicts. A three-layer inspection system of distribution level, constraint level and design availability level is constructed, and experiments and ablation are carried out on the basis of increasing constraint strictness. The results show that the proposed method occupies a Pareto-dominant position in the diversity-harmony plane, and the satisfaction rate of accessibility hard constraints remains above 95% under pressure, which is significantly better than the baseline without explicit modeling constraints. The ablation confirms that relational coding and constraint guidance dominate the gains of harmony and constraint compliance respectively.

## Keywords

Color collocation generation; Discrete diffusion model; Relation-aware coding; Constraint guided denoising; Design usability assessment

**Received: 21 Aug 2026; Revised: 23 Sep 2026; Accepted: 18 Oct 2026; Published: 28 Oct 2026**
**Copyright: ©2026 The Author(s). Published by Digital Intelligence Press Limited. This is an open access article under the CC BY 4.0 license.**

## 1. INTRODUCTION

Color collocation is the visual basis of the design scheme, and its quality directly determines whether the scheme can be read and perceived correctly. However, in the actual design process, color selection is far from a simple externalization of aesthetic intuition, but a complex decision under the simultaneous action of multiple constraints: Functional semantics requires color to bear the responsibility of warning, grading distinction and brand identification; visual harmony requires the relationship between color and lightness to conform to the law of harmony; accessibility requires the contrast between text and background to be no lower than the readable bottom line; style consistency

requires the scheme to fit the established visual tone as a whole. These constraints are not independent of each other but overlap and interact. Strengthening the saturation of the main color of the brand may sacrifice the harmony with the auxiliary color, and lowering the background brightness to highlight the hierarchy may violate the lower limit of contrast [1],[2],[3]. In manual color matching, designers often need to gradually approach feasible solutions through repeated trial and error. The more constraints there are and the narrower the feasible region is, the more significant the efficiency bottleneck becomes. Moreover, the quality of the final scheme is highly dependent on individual experience, which is difficult to reproduce steadily. This reality constitutes the starting point of automatic color collocation generation research: if the algorithm can efficiently explore and output compliant and diverse candidate schemes in the constrained space, the designer's work can be transferred from heavy trial and error to a higher level of judgment and discretion.

The existing automatic generation paths can be roughly classified into three categories, each of which has structural contradictions between constraint controllability and combination diversity [4],[5]. The rule-based method takes the color wheel template and harmonic rule as the core, and Colorgorical tools select combinations from candidate colors through explicit scoring functions [6]. Its advantages are that the harmony degree is controllable and the results can be interpreted. However, the generation space is strictly limited by the template, and the diversity is insufficient, and it is difficult to adapt to complex constraints such as fixed roles and contrast thresholds [7]. Once rules are expanded, they quickly become bloated. The method based on image extracts the color palette from photos or visual manuscripts by clustering, and its diversity is determined by the distribution of the source image [8],[9]. The style is natural but completely subject to the input image, and the constraint cannot be actively adjusted according to the design intention. The method fails when the source image lacks the required color relationship. The text-to-color palette method maps color combinations conditional on semantic labels or descriptions, but its mapping is too deterministic. The same input often converges to similar output, and the combination diversity is limited. Moreover, the semantic-to-color mapping can hardly accommodate accurate numerical constraints such as contrast. The common problem of the three types of methods is that constraints are either hard coded as templates or completely placed outside the generation process, lacking a mechanism that can explicitly carry multiple types of constraints without compressing the diversity of generation.

Recent attempts at deep generative models in color-related tasks suggest that there is a potential breakthrough in this situation. β-CVAE is used to generate color palettes from emotion labels, and the distribution of hidden variables is close to the conditional distribution through conditional prior alignment, which shows better distribution fitting ability than the regular method under unconstrained conditions [10]. The diffusion model is used to colorize grayscale maps, and the pixel is used as the generating unit to gradually denoise, which is outstanding in structure preservation. Together, these works show that generative architectures have the potential to model color distribution, but they also expose an unfilled gap: neither of them uses the quality of the relationship between colors as a guide signal for the generation process. β-CVAE relies on conditional alignment to indirectly affect color, and the relationship constraint is not within the generation target. Image coloring is guided by source graph structure, and the relationship between colors is also outside the generation mechanism. In other words, the existing generation method can generate palette-like combinations but cannot generate relationship-compliant combinations, and the constraint remains in the screening process after output, rather than the intrinsic part of the generation process. This is the position that the technical route of this paper aims to address.

The contribution of this paper can be summarized in three aspects. First, the generation of color

collocation under design constraints is formalized as a sequential decision problem of discrete color combination under multiple constraints. ColorRel-Diff is proposed, which takes discrete color token sequence as the generation unit and uses mask discrete diffusion to explicitly select colors for a functional role for each denoising step. Thus, the generation granularity is shifted from a pixel array to a role-based color combination. Second, the color relation coding module and the constraint guided denoising mechanism are designed. The former encodes the harmonious relation and contrast relation as edge features on the relation graph and aggregates them into relational-perception embedding by graph attention, while the latter introduces the gradient of the relation-consistency score as a sampling guidance term outside the standard denoising objective. The feasibility mask is used to impose inviolable token-level restrictions on hard constraints, so that constraint compliance is endogenously embedded in the generation process rather than dependent on post-processing. Third, a three-layer inspection system of distribution level, constraint level and design availability level is constructed, and the design scenario benchmark with increasing constraint strictness and the component-ablation scheme are matched to provide a reusable paradigm for the evaluation of similar generation tasks. The subsequent sections follow: Section 2 describes the generation architecture, Section 3 describes the constraint coding and injection method, Section 4 establishes the inspection system, Section 5 presents the results and analysis, Section 6 discusses the role and boundaries of the algorithm in the design process, and Section 7 concludes the whole paper.

# 2. DIFFUSION GENERATION ARCHITECTURE OF COLOR RELATIONSHIP PERCEPTION

## 2.1 Overall architecture and discrete diffusion generation objects

ColorRel-Diff represents a set of color collocation of a design scheme as a fixed length token sequence with a sequence length of N (N=5 in this paper, corresponding to the five functional roles of main color, secondary color, emphasis color, background color and text color). Each position is valued at the discrete color vocabulary V. The word list is composed of M color tokens obtained by uniform quantization in CIELAB space (M=1024 in this paper). Let the complete color sequence be $x_0 = (x_0^1, \dots, x_0^N)$, in the forward process, the absorbing state [MASK] replaces the original token position by position according to the scheduling probability. At the $t$ step, the original token is independently retained with probability $\bar{\alpha}_t$ and replaced with [MASK] with probability $1 - \bar{\alpha}_t$. The logistics of marginal distribution of any position are:

$$q(x_t^i \mid x_0^i) = \bar{\alpha}_t \, \delta(x_t^i, x_0^i) + (1 - \bar{\alpha}_t) \, \delta(x_t^i, [\text{MASK}]) \tag{1}$$

Where $x_t^i$ is the token at the $i$-th position in step $t$, $\delta(\cdot,\cdot)$ is the indicator function, and $\bar{\alpha}_t$ is the cumulative reservation probability which decreases monotonically with $t$. This absorbing state structure enables the back sampling to only predict the current mask position without modeling the joint noise for all positions. The reverse denoising network adopts bidirectional Transformer. The input is the current partially visible color sequence, functional role embedding at each position, diffusion step embedding and relational-awareness conditional embedding, and the output is the category distribution of color tokens at each mask position:

$$p_\theta(x_0^i \mid x_t) = \text{softmax}(f_\theta(x_t, i, t, z)) \tag{2}$$

Where $f_\theta$ is the denoising network, $z$ is the relationship-aware conditional embedding, and $\theta$

is the network parameter. During sampling, the decoding sequence with confidence first is adopted: only the batch of positions with the highest predicted confidence is solidified in each step, and the remaining positions are kept as masks to enter the next step [11]. The key to this design is that the generation unit is "color combination" rather than "pixel array", each token in the sequence has its own role semantics, and each step of the denoising process explicitly makes the decision of "selecting a color for a design role", which is essentially different from the pixel-level diffusion of image coloring in terms of problem granularity.

### 2.2 Color relation coding module

The color relationship coding module converts the harmonious relationship and contrast relationship between candidate colors into edge features on the graph structure and then aggregated into the relationship perception embedding of each color taken through graph attention. For any two-color tokens $c_i$ and $c_j$ in the word list, the hue Angle difference $\Delta h_{ij}$, chroma difference $\Delta C_{ij}$ and lightness difference $\Delta L_{ij}$ are calculated in CIELAB space. According to the classical color harmony rules, the degree of the hue Angle difference falling into the four types of intervals of complementarity, analogy, triangle and split complementarity is soft quantified by Gaussian membership degree [12]:

$$m_{ij}^{(k)} = \exp\left(-\frac{\left(\Delta h_{ij} - \mu_k\right)^2}{2\sigma_k^2}\right),\ k \in \{\text{complementary, analogical, triangular, split complementary}\} \quad (3)$$

Where $\mu_k$ is the central hue Angle difference of the $k$-th class rule, and $\sigma_k$ is its allowable bandwidth. The contrast ratio of the two colors is calculated according to the WCAG relative brightness formula [13],[14]:

$$R_{ij} = \frac{L_{\max} + 0.05}{L_{\min} + 0.05} \quad (4)$$

Where $L_{\max}$ and $L_{\min}$ are the larger and smaller luminance of the two-color pairs, and the difference between $R_{ij}$ and threshold $\tau \in \{4.5, 3.0\}$ is mapped to soft satisfaction through sigmoid:

$$s_{ij}^{(\tau)} = \sigma\left(\beta\left(R_{ij} - \tau\right)\right) \quad (5)$$

Where $\beta$ is the steepness parameter. The four-dimensional harmony membership degree and two-dimensional contrast features are spliced, and the edge features are obtained by linear transformation:

$$e_{ij} = W_e\left[m_{ij}^{(1)}, \ldots, m_{ij}^{(4)}, s_{ij}^{(4.5)}, s_{ij}^{(3.0)}\right] + b_e \quad (6)$$

Where $W_e$ and $b_e$ are learnable parameters. The relationship graph takes the word list color as the node and $e_{ij}$ as the edge attribute and adopts the two-layer graph attention network for message transmission [15],[16]. The attention coefficient and aggregation of node $i$ on the $l$-layer to neighbor $j$ are expressed as:

$$\alpha_{ij}^{(l)} = \mathrm{softmax}_j\left(\mathrm{LeakyReLU}(a^\top\left[Wh_i^{(l)} \,\|\, Wh_j^{(l)} \,\|\, W_e e_{ij}\right])\right) \quad (7)$$

$$h_i^{(l+1)} = \sigma\left(\sum\nolimits_{j\in\mathcal{N}(i)} \alpha_{ij}^{(l)} Wh_j^{(l)}\right) \quad (8)$$

Where $h_i^{(l)}$ is the representation of nodes at layer $l$, $a$, $W$ and $W_e$ are learnable parameters, $\|$

represents vector splicing, and $\mathcal{N}(i)$ is the set of neighbors of node $i$. The aggregated node representation is the relationship-aware color embedding $z_i = h_i^{(2)}$. In order to control the cost, this module only calculates offline once on the word list size $M$, and obtains $M \times d$ embedded table ($d = 128$). When generating, it can look up the table according to the token index and does not repeat the calculation with the number of sampling steps. Table 1 shows the main structural parameters of the relational coding module. It can be seen that the number of parameters in this module is of the same order as the scale of the word list rather than the sequence length, so the online computing cost is reduced to a negligible level while maintaining the expression power.

**Table 1. Structural parameters of color relation coding module**

| Term of parameters | Take value | Description |
|---|---|---|
| Vocabulary size M | 1024 | CIELAB quantifies the valid lattice points uniformly |
| Embedding dimension d | 128 | Relational-perception embedding dimension |
| Edge characteristic dimension | 6 | 4-dimensional harmonious membership + 2-dimensional contrast features |
| Graph the number of attention layers | 2 | 4 heads of attention per layer |
| Offline embedding table size | M×d | Generate time lookup table, do not double calculation |

### 2.3 Constraint guided denoising mechanism

The constraining guidance denoising mechanism is the core innovation of this section, which includes two complementary paths: soft guidance and hard mask. The soft guidance path introduces the relationship consistency score as a differentiable guidance signal. Given the determined color set and the candidate distribution of the position to be predicted in the current sequence, the relationship consistency score is defined as the expectation of the sum of the relationship scores of all colors in the sequence [17],[18]:

$$S(x) = \mathbb{E}\left[\sum_{i<j} g_\phi\left(z_{x^i}, z_{x^j}, e_{x^i x^j}\right)\right] \tag{9}$$

Where $g_\phi$ is a small scoring head, the two-color relationship perception embedding and edge features are mapped into scalar relationship scores, and $\phi$ is the scoring head parameter. The output logits of the denoising network at step $t$ is $\ell_t$, and the sampling distribution after guidance is composed of the original distribution and the scoring gradient correction term [19],[20]:

$$\tilde{p}_\theta\left(x_0^i \mid x_t\right) = \mathrm{softmax}\left(\ell_t^i + \lambda \nabla_{\ell_t^i} S(x)\right) \tag{10}$$

Where $\lambda$ is the guidance strength, and $\nabla_{\ell_t^i} S(x)$ is the gradient of the relationship consistency score to the logits at position $i$, whose effect is to shift the sampling trajectory toward the region with higher color relationship score. Since the score can be calculated based on the pair combination of the offline embedding table and the current sequence, the bootstrap does not need to backpropagation the generation process but only needs to score the candidate tokens one by one and do SoftMax reweighting during sampling. Hard mask path processing cannot violate the constraint: For fixed location constraint,

the corresponding location token distribution will be reduced to the point distribution on the brand color token:

$$p\left(x_0^i = c_{\text{brand}}\right) = 1,\ p\left(x_0^i \neq c_{\text{brand}}\right) = 0 \tag{11}$$

For the background-text contrast constraint, a feasibility mask is constructed before each sampling step, and the logits of candidate tokens that violate the threshold are set to negative infinity:

$$\ell_t^{i,j} \leftarrow -\infty \text{ If } R_{c_j, c_{\text{bg}}} < \tau \tag{12}$$

Where $c_j$ is the candidate token of position $i$, and $c_{\text{bg}}$ is the determined background color. When the hard constraint set cannot be satisfied, the system relates the soft constraint step by step according to the preset priority and outputs the minimum violation scheme, with the conflict location hint. Table 2 lists the key hyperparameters and their action positions of the guidance and mask mechanisms. It can be seen that the soft and hard paths act on different links in the sampling process: soft guidance modulates the shape of the distribution, and hard mask truncates the infeasible support set. After the superposition of the two paths, sampling still maintains the ability to explore the color combination in the feasible region.

**Table 2. Key hyperparameters and action positions of constraint-guided denoising mechanism**

| The Mechanics | Key parameters | Position of action | Nature of action |
|---|---|---|---|
| Soft guidance | The guidance strength λ=0.5 | The sampling distribution is reweighted at each step | Shape of modulation distribution |
| Soft guidance | The rating head has a hidden dimension of 64 | The relational edge represents a → scalar | Differentiable score |
| Hard mask | Contrast threshold 4.5/3.0 | Pairwise token logits | Truncate the set of infeasible supports |
| Hard mask | The number of fixed positions is ≤2 | Corresponding location point distribution | Inviolable constraint |
| Conflict handling | Priority level 3 | When the constraint is not satisfied | Minimum violation output |

## 2.4 Training strategy and course learning

The training is divided into two stages to alleviate the problem of scarce annotated data with constraints. In the first stage, the basic denoising network and relation scoring head are pre-trained on large-scale unconstrained color palette data, and the training target is weighted cross-entropy of discrete diffusion of standard masks:

$$\mathcal{L}_{\text{pre}} = \mathbb{E}_{t,x_0}\left[\sum_{i \in \mathcal{M}_t} w(t)\left(-\log p_\theta\left(x_0^i \mid x_t\right)\right)\right] \tag{13}$$

Where $\mathcal{M}_t$ is the set of masked positions at the $t$-th step, and $w(t)$ is the loss weight scheduled according to the diffusion step, which is used to balance the gradient contribution at different noise levels. In this stage, the model learns the basic distribution of "what color combination looks like the real design collocation", and the embedding table of the relational coding module participates in the training synchronously in this stage. In the second stage, the guidance related parameters are fine-tuned

on the constrained annotation subset: the offline embedding table of the relational coding module and the bottom layer of the denoising network are frozen, and only the conditional modulation parameters of the scoring head and the top layer of the denoising network are updated. In addition to the original negative log-likelihood, constraints are added to the training target to meet the auxiliary loss:

$$\mathcal{L}_{\text{ft}} = \mathcal{L}_{\text{pre}} + \gamma\, \mathbb{E}[\sum_{k} \max(0, \tau_k - R_k(x))] \tag{14}$$

Where $\gamma$ is the auxiliary loss weight, $\tau_k$ and $R_k(x)$ are the threshold value of the $k$-th constraint and the corresponding measurement value of the current sequence respectively, and the second term imposes punishment on the samples violating the constraint. The scheduling of course learning is reflected in the second stage: firstly, samples containing only harmony constraints are trained, and then accessibility constraints and role fixed constraints are gradually introduced. The proportion of samples of each level of constraints increases according to the preset course, so as to avoid underfitting of the model due to gradient spareness caused by hard constraint masks in the early training stage. By comparing the configuration differences between the two training stages in Table 3, it can be seen that the number of trainable parameters in the fine-tuning stage is significantly compressed, which not only reduces the risk of overfitting on small-scale constraint data, but also retains the knowledge of the basic distribution.

**Table 3. Comparison of two-stage training configurations**

| A configuration item | Stage 1 (pre-training) | Stage 2 (Constraint fine-tuning) |
|---|---|---|
| The training data | Large scale unconstrained color palette | Annotated subset with constraints |
| Trainable part | Denoising network + scoring head + embedding table | Scoring head + top conditioning modulation |
| The main losses | Mask position weighted cross entropy | The cross-entropy + constraint satisfies the auxiliary loss |
| Course scheduling | No | Harmony → accessibility → Fixed roles |
| Rate of learning | 1e-4 | 2e-5 |

## 2.5 Positioning of differences from existing generation architectures

The difference between this architecture and the existing generation path can be located from three dimensions: generation unit, condition source and guidance signal. In the generation unit, the $\beta$-CVAE color palette generation method outputs color parameters through the decoder with continuous hidden variables. The generated object is continuous color vector, which needs post-processing discretization to drop into the finite color vocabulary, while ColorRel-Diff directly does mask diffusion on the discrete token sequence. The generated results are naturally valid colors in the vocabulary, and each step of the decision can be traced; The image coloring class diffusion method takes pixels as the generation unit, and its goal is to assign color to a given grayscale map. The output is strongly constrained by the structure of the source map, which is different from the problem setting of the proposed architecture which generates a set of role-oriented color schemes from scratch. In terms of condition sources, the existing methods mostly take semantic categories, emotional labels or reference images as conditions. The conditional core of this architecture is the color relationship diagram itself. Functional roles, accessibility thresholds and style reference are injected through relation coding or mask, and the

conditions are in the same color semantic space as the generating unit. For the guidance signal, $\beta$-CVAE relies on the conditional prior alignment to make the distribution of hidden variables close to the conditional distribution, and the image coloring depends on the structure of the source map. Both $\beta$-CVAE and $\beta$-CVAE do not take "the quality of the relationship between colors" as the explicit guidance term of the sampling period. The relationship consistency score gradient of ColorRel-Diff directly acts on the sampling distribution, so that the sampling trajectory shifts to the region of high relationship quality while meeting the data distribution, which is the fundamental difference between this architecture and the above two methods. The above differences determine that the three are in different positions in the design process: the coloring method serves the color matching adjustment of the existing visual draft, $\beta$-CVAE serves the mapping of emotional labels to the color palette, and ColorRel-Diff serves the generation of color schemes with functional roles and hard constraints. The three cannot be replaced with each other.

# 3. FORMAL CODING AND INJECTION METHODS OF DESIGN CONSTRAINTS

## 3.1 Positional Prior coding and injection of functional semantic constraints

The functional role in the design scene is the structural prior of the color scheme: the main color bears the brand identity, the secondary color supports the level transition, the emphasis color is used for local guidance, and the background color and the text color jointly determine the reading substrate [21],[22]. ColorRel-Diff encodes this prior into position role embedding, which is spliced with token embedding and diffusion step embedding in the input layer of denoising network, and modulates the role condition of features in the middle layer through adaptive layer normalization. Let the discrete role label of the $i$-th position be $r_i \in \{1, \dots, 5\}$, the corresponding learnable role embedding table is $E_r$, then on the input side, the input at position $i$ at step $t$ is expressed as:

$$u_i^{(0)} = E_x(x_t^i) + E_r(r_i) + E_t(t) + P_i \tag{15}$$

Where $E_x$ is the color token embedding, $E_t$ is the diffusion step embedding and $P_i$ is the location embedding. Role modulation at the middle layer adopts adaptive layer normalization, so that role information is not limited to input splicing but continuously affects feature distribution at each layer:

$$\mathrm{AdaLN}(h_i \mid r_i) = \gamma(r_i) \cdot \frac{h_i - \mu(h_i)}{\sigma(h_i)} + \beta(r_i) \tag{16}$$

Where $\mu(\cdot)$ and $\sigma(\cdot)$ are the mean and standard deviation of the feature dimension, and $\gamma(r_i)$ and $\beta(r_i)$ are the dimensional scaling and offset parameters generated by the two-layer perceptrons of role embedding. In this injection method, the character tendencies such as "main color should be more saturated and background color should be less saturated" do not need to be hard-coded in the form of rules but are reflected in the modulation parameters driven by data. On this basis, part of the constraint mode is realized through the role mask: when the user specified brand color is fixed at the position of the main color and other roles are generated freely, the system degrades the token distribution of the position of the main color to the brand color point distribution, and only performs diffusion sampling for the other four positions, and the sampling target can be written as:

$$p_\theta\left(x_0^{\setminus k} \mid x_t^{\setminus k}, x_0^k = c_{\text{brand}}, z, r\right) = \prod_{i \neq k} p_\theta\left(x_0^i \mid x_t\right) \tag{17}$$

Where $k$ is the fixed brand color position, $x_0^{\setminus k}$ represents the color sequence of other positions except $k$, and $z$ is the relational perception conditional embedding. The engineering significance of this model is that in design practice, brand color is often locked by the existing visual norms, and the number of characters that can be generated freely is usually three to four. Table 4 shows the configuration of the three types of partial constraint modes and their compression degree to the generation degree of freedom. It can be seen that for each increase in the number of fixed positions, the number of feasible combinations decreases according to the scale of the word list, while the role conditional modulation ensures that the remaining positions can still generate schemes coordinated with the fixed color.

**Table 4. Constraint patterns and generating degrees of freedom in functional semantics**

| Mode of constraint | Fixed position | Number of free positions | Feasible combination magnitude | Typical scenario |
|---|---|---|---|---|
| Full freedom | No | 5 | $M^5$ | Exploration of concepts |
| Brand lock-in | Main color | 4 | $M^4$ | Brand extension design |
| Brand + substrate locking | Main color, background color | 3 | $M^3$ | Poster/interface color matching |

### 3.2 Feasibility mask coding of accessibility constraints

The core of accessibility constraint is the lower limit of contrast between text color and background color, and its mathematical form is the feasibility constraint between pairs of tokens. The injection position is logits truncation on the sampling side, rather than the loss penalty on the training side. According to WCAG 2.1, the contrast of two colors is defined by the relative brightness ratio, and the text-background combination must meet the requirements of no less than 4.5 for regular text and no less than 3.0 for large text. Set the background color position as $b$ and the text color position as $f$. When sampling at step $t$, construct feasibility indication for the text position candidate token $c_j$:

$$\mathbb{1}_{\text{acc}}(c_j, c_{\text{bg}}) = \begin{cases} 1, & R(c_j, c_{\text{bg}}) \geq \tau_{\text{txt}} \\ 0, & \text{otherwise} \end{cases} \tag{18}$$

Where $R(\cdot,\cdot)$is the contrast ratio defined in Section 2, and $\tau_{\text{txt}}$ is the text file threshold. The injection method is to impose a negative infinite penalty on candidates who do not meet the feasibility before SoftMax:

$$\tilde{\ell}_t^{f,j} = \ell_t^{f,j} + \log \mathbb{1}_{\text{acc}}\left(c_j, c_{\text{bg}}\right) \tag{19}$$

Since the determination sequence of background color and text color is not fixed in confidence priority decoding, the application timing of the mask needs to be linked with the decoding sequence: if the background color is determined before the text color, the above mask is applied to the text position. If the text color is determined first, a symmetric mask is applied to the background position. When both are solidified in the same batch, a joint feasibility check is applied and only token pair combinations that satisfy the constraints are retained. This design makes accessibility constraint an inviolable hard

boundary at the sampling level, and the results generated can ensure the bottom line of readability without post-processing and filtering. Furthermore, accessibility constraints are not limited to text-background pairs: when the emphasis color is used to carry key information ICONS, the contrast between the emphasis color and the background color can also be included in the mask set, and the feasibility indication is extended to the combination of multiple paired constraints:

$$\mathbb{1}_{\text{acc}}^{\text{multi}}(c_j) = \prod_{(p,q)\in\mathcal{P}} \mathbb{1}_{\text{acc}}\left(c_j^{(p)}, c^{(q)}\right) \tag{20}$$

Where $\mathcal{P}$ is the set of role pairs that need to meet the contrast requirements, and $c_j^{(p)}$ represents the color when candidate $c_j$ is placed in role $p$. This extension does not change the structure of the mask mechanism and only increases the number of pairs of feasibility checks at each step. Table 5 lists the key configurations of accessibility masks. It can be seen that threshold gear and role pair set are two independent adjustable dimensions. The former determines the strictness, while the latter determines the constraint coverage, which jointly describes the compliance level of the scheme in the readability dimension.

**Table 5. Key configurations of accessibility feasibility masks**

| A configuration item | Value/range | Effect of action |
|---|---|---|
| Text contrast threshold $\tau_{txt}$ | 4.5 (regular) /3.0 (large text) | Determine the text-background compliance bottom line |
| Graphic contrast threshold $\tau_{ui}$ | 3.0 | Emphasize color-background icon readability |
| Constrain the set of role pairs P | {text, background},{emphasis, background} | Determine the constraint coverage |
| The mask applies timing | Linkage with decoding sequence | Ensure that pairwise checks are valid |
| Multiple pairs of constraint combinations | Conjunctive | Sampling can be done only when all are met |

### 3.3 Soft alignment coding with style tendency constraints

The style tendency constraint is different from the previous two types of constraints in mathematical form: it does not appear as a hard boundary of "must be satisfied", but as a soft target of "as close as possible", whose carrier is a single sheet or a small reference palette. The key to encoding the reference palette as a soft constraint is to construct a similarity measure that is insensitive to the set order, since the reference palette and the generation scheme do not necessarily correspond in role order. In this paper, the palette difference measure based on Hungarian matching is adopted: let the reference palette be $Q = \{q_1, \ldots, q_K\}$, the generation scheme is $X = \{x^1, \ldots, x^N\}$, the two-color distance is defined as the weighted Euclidean distance in CIELAB space:

$$d(q_m, x^i) = \sqrt{w_L(L_m - L_i)^2 + w_a(a_m - a_i)^2 + w_b(b_m - b_i)^2} \tag{21}$$

Where $w_L, w_a, w_b$ are the three-channel weights, which are used to regulate the relative importance of lightness and chroma in style similarity. The palette difference is defined as the minimum total distance under optimal one-to-one matching:

$$D(Q,X) = \min_{\pi\in\Pi} \sum_{m=1}^{K} d\left(q_m, x^{\pi(m)}\right) \tag{22}$$

Where $\Pi$ is the set of one-to-one mappings from the reference color to the generated color. When $K < N$, colors in the generated scheme are allowed not to participate in matching. Style alignment loss is obtained from this difference by monotone transformation:

$$\mathcal{L}_{\text{style}} = \log(1 + D(Q,X)) \tag{23}$$

Its injection position is on the guidance score side, and the weighted combination with the relationship consistency score in Section 2 is the total guidance target:

$$S_{\text{total}}(X) = S_{\text{rel}}(X) - \eta\,\mathcal{L}_{\text{style}}(Q,X) \tag{24}$$

Where $S_{\text{rel}}$ is the relationship consistency score, $\eta$ is the style weight, and the negative sign indicates that the smaller the style difference is, the higher the guide score is. During sampling, the logits gradient is calculated for the total guidance target, so that the sampling trajectory is biased towards the regions with high relationship quality and high style similarity [23]. The advantage of this coding method is that the number of colors in the reference palette and the number of characters in the generation scheme need not be the same, and the Hungarian matching automatically deals with the case of unequal number and order; At the same time, style constraints enter the guide item in a soft form. When the style reference conflicts with the direction of functional semantics or hard constraints of accessibility, hard constraints take precedence, and style alignment automatically retreats without explicit conflict handling. Table 6 shows the key parameters of style soft alignment and their influence on guiding behavior. It can be seen that the weight $\eta$ and the distance weight $w_L$ are the two main regulating buttons. The former controls the relative discourse power of style constraints, while the latter controls whether the style similarity focuses on the lightness structure or the hue structure.

**Table 6. Key parameters of style soft alignment and bootstrapping influence**

| Term of parameters | Take value |
|---|---|
| Style weight $\eta$ | 0.3 |
| Lightness weight $w_L$ | 1.0 |
| The chromaticity weights $w_a$, $w_b$ | 0.5 |
| Reference chromatic number $K$ | 1-5 |
| Mode of matching | The Hungarian optimal matching |

### 3.4 Constraint conflict detection and grading softening

The three types of constraints can be satisfied at the same time in most scenarios, but there may be structural conflicts between the hard constraints: if the contrast between the fixed brand color and the white background is insufficient, "the brand color is fixed at the main color" and "the main color-background contrast reaches the standard" cannot be established simultaneously. For example, when the background color is locked to dark and the text color is locked to medium gray, the text-background contrast constraint cannot be met. Conflict detection is performed before sampling begins check

whether the intersection of the feasible region is empty for the hard constraint set, check whether the fixed color pair has violated the threshold for the pair contrast constraint, and check whether it contradicts the role pair constraint for the fixed position constraint [24]. After a conflict is detected, the system does not fail directly but softens the constraints gratingly according to the preset priority. The priorities from high to low are accessibility bottom line, brand color fixed, character integrity, style orientation, based on the fact that readability relates to the functional usability of the solution, brand color relates to asset consistency, the two cannot be easily compromised, while character integrity and style orientation are adjustable aesthetic goals. The softening strategy introduces slack variables to each level of constraints, transforms hard constraints into soft constraints with penalties, and solves the minimum violation scheme:

$$\min_X \sum_{k\in\mathcal{C}_{\text{soft}}} \omega_k \ \delta_k(X) \text{ s.t. } \delta_k(X) \geq 0, \ \sum_k \delta_k(X) \leq \epsilon \tag{25}$$

Where $\mathcal{C}_{\text{soft}}$ is the softened constraint set, $\delta_k(X)$ is the amount of violation of the $k$-th constraint, $\omega_k$ is the penalty weight set according to priority, and $\epsilon$ is the total violation budget. The solution results are accompanied by a conflict prompt indicating the softened constraint, the amount of violation, and the suggested direction of manual adjustment (e.g., "It is recommended to replace the light variant of the brand color to restore contrast"). This mechanism allows the system to output a usable solution in the face of unsatisfiable constraints, rather than terminating the process with an empty result, while returning the final discretion of conflict decisions to the designer. Table 7 shows the conflict types and corresponding softening disposal.

**Table 7. Types of constraint conflicts and graded softening disposition**

| Type of conflict | Condition of inspection | Softening disposal | Whether to prompt |
|---|---|---|---|
| Brand color and contrast conflict | Fixed color pair R<τ | Retain the brand color, soften the contrast and inform with a prompt | Yes |
| Double fixed color contrast conflict | Background/text is fixed and R<τ | Soften text color fixed, allowing fine tuning | Yes |
| Role integrity and fixed conflict | Fixed number of positions > Number of roles | Softening character integrity | Yes |
| Style and hard constraints conflict | The style-optimal solution violates the hard constraint | Style automatically retreats and does not account for budget violations | No |

To sum up, the three types of constraints are injected into the input side, sampling side and scoring side respectively in three different mathematical forms: position priority, feasibility mask and soft alignment loss, forming a grading constraint system. Conflict detection and grading softening ensure that the system can still produce the minimum violation scheme when the constraints cannot be satisfied simultaneously.

# 4. EXPERIMENT DESIGN AND TEST SYSTEM

## 4.1 Layer 1: Generate distribution quality test

The goal of layer 1 is to answer, "whether the generated color combination is like a reasonable

design palette distribution as a whole". The test object is the set of schemes output by the method rather than a single scheme, so the indicators are defined at the set level. Color diversity describes the coverage uniformity of the scheme on the hue ring: project the color of each character of each scheme to the CIELAB hue Angle, make a weighted histogram of all colors of all schemes, and measure the coverage deviation by the histogram and the uniformly distributed Jensen-Shannon divergence:

$$\mathrm{Div} = 1 - \mathrm{JSD}(p_h \parallel u) \tag{26}$$

Where $p_h$ is the weighted histogram of hue Angle, $u$ is, the uniform distribution, and the closer the value is to 1, the more uniform the hue coverage is. The combination novelty describes the non-memory of the generation scheme relative to the training set: find the nearest neighbor of each generation scheme in the training set, use the palette difference measure (defined in Section 3) as the distance, and take the mean and low quantile of the nearest neighbor distance of all schemes as the novelty index:

$$\mathrm{Nov} = \mathbb{E}_{X\sim p_\theta}\left[\min_{Q\in\mathcal{D}_{\text{train}}} D(Q,X)\right] \tag{27}$$

Where $\mathcal{D}_{\text{train}}$ is the training set palette set. This indicator should be interpreted jointly with diversity: only high novelty may mean that the generated results deviate from a reasonable distribution, and only high diversity may mean random noise, and the two must be in a reasonable range at the same time to be convincing. According to the four types of harmony membership defined in Section 2, the mean value of all color pairs in the scheme is calculated, and the contrast soft satisfaction is superimposed:

$$\mathrm{Harm}(X) = \frac{1}{\binom{N}{2}}\sum_{i<j}\left(\max_k m_{ij}^{(k)} + \rho\, s_{ij}^{(4.5)}\right) \tag{28}$$

Where $\rho$ is the weight of the contrast term, which is used to balance the relative contribution of the harmony rule and readability in the rating. The score does not depend on the model in this paper and is therefore comparable to all baseline methods. The baseline selection covers three types of representative paths: Colorgorical rule generation represents the deterministic method based on color harmonic template, K-means image extraction palette represents the data-driven extraction method, and the adaptive version of $\rho$-CVAE emotion palette generation represents the continuous hidden variable generation method. In terms of generation mechanism, the three routes belong to rule, extraction and generation, which constitute a differentiated comparison with the method in this paper. To control comparative fairness, all generative methods use the same size of sampling budget, and the rules and extraction methods run according to their default configurations. **Table 8** shows the construction elements and baseline applicability of each index at layer 1.

**Table 8. Construction and applicability of quality indicators for Layer I generation distribution**

| Indicators | Foundation of structure | Direction of value | Baseline of application |
|---|---|---|---|
| Diversity Div | Hue histogram with uniformly distributed JSD | The closer it gets to 1, the better | All |
| Novelty Nov | The nearest neighbor palette distance of the training set means | Moderately high is appropriate | All |
| Harmony degree Harm | Harmonious membership + contrast soft satisfaction | The higher the better | All |
| Budget for sampling | Number of schemes generated per method | Unification | Generative method |

### 4.2 Layer 2: Constraint satisfaction rate test

Layer 2 tests the actual effectiveness of the constraint coding mechanism in Section 3. Its core design is to construct three test sets with increasing constraint strictness, so that the method can expose differences under different constraint pressures. Test set T1 only contains harmony constraints, that is, the color pairs in the scheme are required to meet at least one type of harmony rules; On the basis of T1, test set T2 adds hard accessibility constraints, requiring text-background contrast to be no less than 4.5 and emphasis on-background to be no less than 3.0. Test set T3 adds some role fixing on the basis of T2, requiring the main color position to be fixed as a given brand color and the other four characters to be generated freely. The number of constraints and the volume of the feasible region of the three test sets are successively narrowed to form the pressure gradient. The hard constraint satisfaction rate is defined as the passing proportion of the scheme on all hard constraints:

$$\mathrm{Sat}_{\mathrm{hard}} = \frac{1}{|\mathcal{X}|} \sum_{X \in \mathcal{X}} \prod_{k \in \mathcal{C}_{\mathrm{hard}}} \mathbb{1}[\delta_k(X) = 0] \tag{29}$$

Where $\mathcal{X}$ is the set of generating schemes, $\mathcal{C}_{\mathrm{hard}}$ is the set of hard constraints, $\delta_k(X)$ is the violation of the $k$-th constraint, and the indicator function takes the value of 1 when the constraint is fully satisfied. The degree of soft constraint violation takes the meaning of the normalized violation amount of each soft constraint:

$$\mathrm{Vio}_{\mathrm{soft}} = \frac{1}{|\mathcal{X}|} \sum_{X \in \mathcal{X}} \sum_{k \in \mathcal{C}_{\mathrm{soft}}} \omega_k \frac{\delta_k(X)}{\delta_k^{\max}} \tag{30}$$

Where $\omega_k$ is the $k$-th soft constraint weight, and $\delta_k^{\max}$ is its normalization benchmark. For the difference of satisfaction rate between groups, McNemar test is used to compare the pass/fail pairing results of the proposed method and each baseline on the same test sample. The test statistics are as follows:

$$\chi^2 = \frac{(b-c)^2}{b+c} \tag{31}$$

Where $b$ and $c$ are the logarithm of the scheme that passed the method in this paper but did not pass the baseline, and the scheme that passed the baseline but did not pass the plan in this paper respectively, which obey the chi-square distribution with 1 degree of freedom under the null hypothesis. For the confidence interval of satisfaction rate and violation degree, the bootstrap method is used to resample: the scheme set is put back and extracted B times (B=1000 in this paper), and the index is recalculated each time, and the quantiles of 2.5% and 97.5% are taken to form the 95% confidence interval. The bootstrap method does not rely on the normality assumption of the index and is suitable for a bounded proportion index such as the satisfaction rate. Table 9 shows the corresponding relationship between constraint configuration and test methods of the three test sets. It can be seen that the number of constraints from T1 to T3 increases and the feasible region Narrows in turn. The difference between methods is expected to be the most significant in T2 and T3, because all generative methods can achieve high satisfaction rate through post-processing filtering under T1. However, the difference in satisfaction rates under conditions T2 and T3 directly reflects whether the constraints are endogenously encoded into the generation process.

**Table 9. Configuration of the three-constraint strictness increasing test sets in Layer 2**

| Set of tests | Composition of constraints | Region of feasibility | Main test |
|---|---|---|---|
| T1 | Constraint of harmony | wide | Harmonious satisfaction rate, McNemar |
| T2 | Harmony + accessibility | medium | Hard constraint satisfaction rate, self-help CI |
| T3 | Harmony + accessibility + role fixation | narrow | Satisfaction rate, degree of soft violation |

### 4.3 Layer 3: Design usability agent evaluation

The problem faced by layer 3 is that satisfying constraints and reasonable distribution do not automatically mean that the scheme is usable by design. Direct subjective human assessment involves ethical approval and questionnaire design, which is not adopted in this paper. The calibration set of the agent evaluation is composed of senior students at the School of Design (n=12), who do a quick ranking of the generated schemes within the framework of course assignments. The ranking task only involves the relative order of the schemes in the three dimensions of "scene suitability", "visual hierarchy clarity" and "overall preference" and does not collect sensitive personal information or make individual level inference. Therefore, it does not constitute human subject research requiring ethical approval. The Bradley-Terry model is used to analyze the ranking data, and the probability that alternative $i$ is better than alternative $j$ is modeled as follows [25],[26],[27]:

$$P(i \succ j) = \frac{\exp(\theta_i)}{\exp(\theta_i) + \exp(\theta_j)} \tag{32}$$

Where $\theta_i$ is the potential usability score of schemes $i$, and the average score of each method scheme obtained by maximum likelihood estimation is used as the proxy value of design usability of this method. The validity of the proxy index should be verified by correlation analysis: Spearman rank correlation was made between the automatic index value of each method at layer 1 and layer 2 and the Bradley-Terry score at layer 3. If an automatic index was significantly correlated with the proxy score, it indicated that the index could represent the design judgment to a certain extent. If the correlation is not significant, it suggests that the indicator only describes the distribution side unrelated to design

availability. This verification step is the key to the closure of the three-layer system, which makes the automatic indicators of layer 1 and layer 2 not assumed to be sufficient representatives of availability by default, but to test their representative ability with data. Table 10 presents the components of the three-agent evaluation at the layer, and it can be seen that the design intention of this layer is not to replace large-scale subjective research with 12-person ranking, but to establish the corresponding relationship between automatic indicators and design judgment with the minimum scale calibration set, thus providing a reusable verification paradigm for subsequent research.

**Table 10. Building blocks of Tier 3 design usability proxy evaluation**

| The elements | Configuration | Description |
|---|---|---|
| Calibration set size | n=12 | Senior student at the School of Design |
| Dimensions of assessment | Scene fit/hierarchy clarity/overall preference | Quick sort, no scale |
| Model of score | Bradley-Terry | The potential score is estimated by maximum likelihood |
| Test of validity | Spearman's rank correlation | Automatic indicator vs proxy score |
| Nature of data | Program sequencing within coursework | No sensitive information is collected |

### 4.4 Ablation experiment design

The purpose of the ablation experiment is to attribute the overall performance of the method in this paper to specific components, so the stripping scheme should correspond one-to-one to the architectural components in Section 2 and keep the rest operational after stripping. Four ablation configuration were set: complete model, removal of relational coding module (learnable free token embedding was used to replace relational sensing embedding, and the rest of the denoising network remained unchanged), removal of constraint guidance term (hard mask was retained, only relationship consistency score gradient guidance was removed), and removal of course learning (two stages were merged into a single stage for training on mixed data). All ablation configurations were evaluated on layer 1 and Layer 2 indices, and their degradation amplitude was compared with that of the full model:

$$\Delta_{\text{abl}} = \frac{M_{\text{full}} - M_{\text{abl}}}{M_{\text{full}}} \times 100\% \tag{33}$$

Where $M$ is the index to be compared (diversity, harmony, hard constraint satisfaction rate, etc.), $M_{\text{full}}$ is the value of the complete model, and $M_{\text{abl}}$ is the value of the ablation configuration. The interpretation of degradation degree should be combined with the function of components: if the removal of relationship coding module mainly leads to the decrease of harmony degree but little change of constraint satisfaction rate, it indicates that the function of this module focuses on relationship quality rather than constraint compliance; If the removal of constraint guidance item mainly leads to the increase of soft constraint violation while the satisfaction rate of hard constraint remains unchanged, it indicates that there is a division of labor between soft guidance and hard mask in function. Ablation of course learning requires additional reporting of training convergence curves to observe whether single-stage training underfits on hard constrained samples. Table 11 lists the indicator dimensions affected by the four ablation configurations and expectations.

**Table 11. Ablation experiment configuration and expected degradation pattern**

| Configuration of ablation | Remove an object | Expected main affected indicators | Significance of comparison |
|---|---|---|---|
| The complete model | No | - | Benchmark |
| w/o relational coding | Relationship-aware embedding | Harmony, novelty | Relationship quality attribution |
| w/o boot entry | Scoring gradient guidance | Degree of soft constraint violation | Attribution of hard and soft division of labor |
| w/o course learning | Two stage scheduling | Convergence stability, T3 satisfaction rate | Training strategic attribution |

To sum up, the three-layer test system describes different aspects of the results generated by the indicators of the collection layer, the satisfaction rate of the constraint layer and the ranking score of the agent layer respectively.

# 5. RESULTS AND ANALYSIS

## 5.1 Quality of distribution: dominant position on the Pareto front

The core of layer 1 results is not the highest value of a certain index, but the joint occupancy in the plane of diversity and harmony. Figure 1 presents the joint distribution of each method in the form of scatter points on diversity (horizontal axis) and harmony (vertical axis), with each point being the recalcination result of a sampling batch, and the dotted line being the Pareto frontier envelope of each method. It can be seen that the point group of the rule generation method is concentrated in the narrow upper right corner area with high harmony and low diversity, and its leading harmony degree comes from the direct constraint of the harmonic template. However, the hue coverage is seriously biased towards the few Angle difference interval allowed by the template, and diversity collapse is the inevitable cost of the template prior. $K$-means image extraction method is on the contrary, the point group expands along the bottom right direction, with high diversity but significant variance of harmony degree. The reason is that the extraction method faithfully replicates the color distribution of the source map, and the harmony of the source map itself is not controlled, and the large variance is the direct consequence of the lack of relationship constraints. The position of the $\beta$-CVAE adaptation is close to ColorRel-Diff under unconstrained conditions, which indicates that the continuous hidden variable generation path is not inferior in fitting the edge distribution of the color palette. However, the method in this paper has a higher median of harmony at the same level of diversity, and the point group is closer to the frontier envelope, indicating that the role of relational coding and guidance term is not to elevate a single index, but to push the generated results to the region where both indexes are better at the same time. In the interpretation of Figure 1, it should be noted that the Pareto dominance relationship is at the aggregate level, and the overlap of individual batches does not constitute the negation of the dominance relationship.

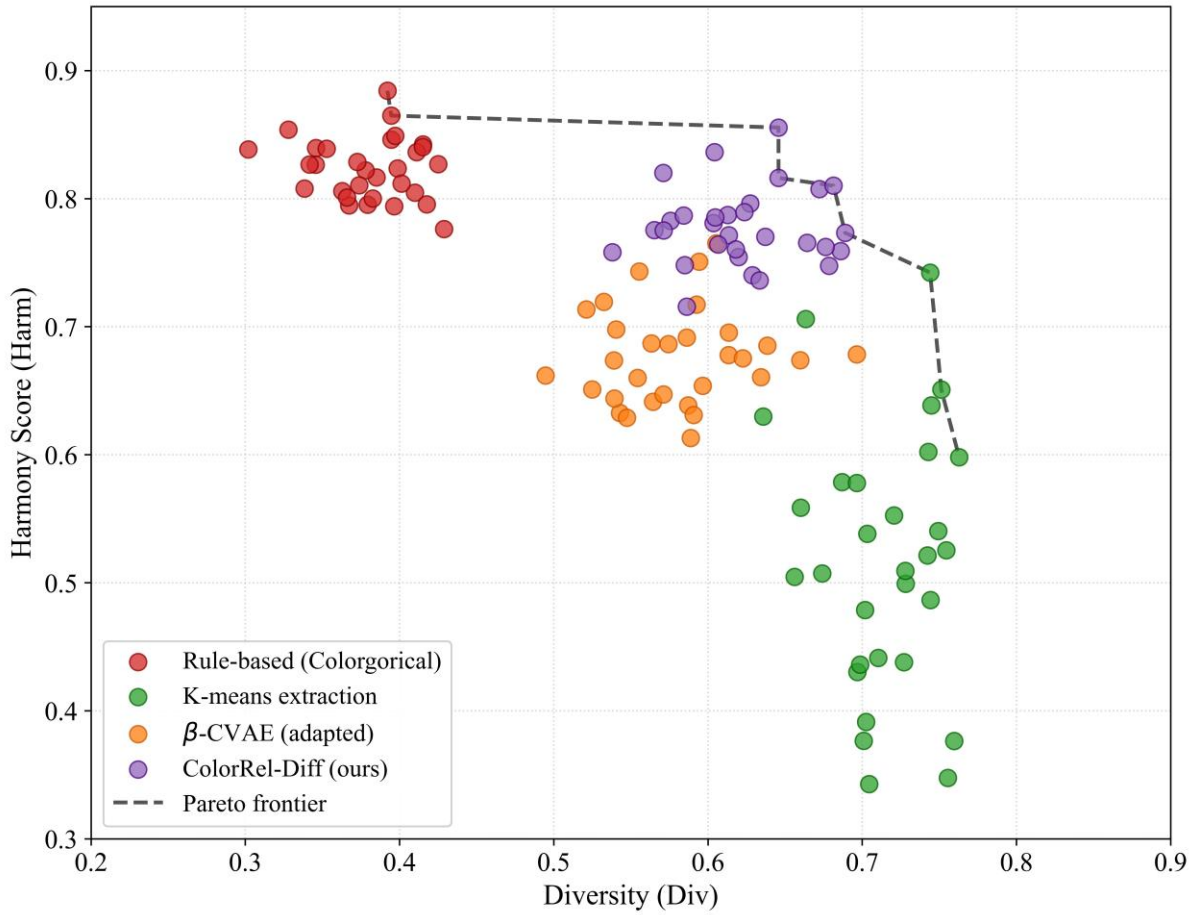



**Figure 1. Joint distribution and Pareto frontier envelope of each method in the diversity-harmony plane**

The novelty indicator provides a complementary constraint to the above judgment. If the high degree of harmony of the method in this paper comes from the memory of the training set, the novelty should be significantly low. β-CVAE is close to the method in this paper in novelty, but its distribution shift problem in the generation of role conditions is fully exposed at layer two, which is left to be discussed in the next section. Based on the results of layer 1, the causal explanation can be given as follows: The high harmony degree of the rule method is the product of "too strong constraints", and the high diversity of the extraction method is the product of "absence of constraints". The joint dominance of the method in this paper comes from injecting relational constraints into the generation process in the form of soft guidance, so that diversity does not need to be at the cost of harmony. This mechanism has been effective under the condition of no constraints. However, the magnitude of advantage is limited by the competitiveness of strong baselines such as $\beta$-CVAE in the fitting of marginal distributions.

### 5.2 Constraint satisfaction: Advantages are amplified as constraints are tightened

The second layer result is the most distinctive part of this section, and its narrative logic is "how the gap between methods changes as the constraint stringency increases." On T1 with harmony constraints only, there is no significant difference in the satisfaction rates of hard constraints among the generative methods, and the paired pass/fail counts of McNemar test are close to symmetry, which is in line with expectations: The harmony constraint is relatively loose, and each method can achieve a high

satisfaction rate by filtering after sampling or biasing the distribution itself. The real differentiation occurs in T2: after introducing the hard constraints of text-background and emphasis-background accessibility, the satisfaction rate of methods without explicit modeling constraints decreases significantly. The reason is that the sampling process of these methods is not aware of the pairwise contrast relationship, and high-contrast combinations only account for a minority in their output distribution. A large number of schemes will be discarded at the same time by filtering post-processing, and the effective output rate will decrease accordingly. The feasibility mask of ColorRel-Diff truncates the infeasible support set on the sampling side, and the satisfaction rate is maintained at a high level, and the effective output rate is not lost due to filtering. Figure 2 presents the hard constraint satisfaction rates of each method on the three test sets in the form of grouped columns. The column height is the point estimate, and the error line is the 95% confidence interval of the bootstrap method. It can be intuitively seen that the column height on T1 is close, and there is an obvious gap between the proposed method and other methods on T2 and T3, which expands with the increase of constraint strictness. This is the inevitable distinction between "constrained endogenous coding" and "constrained post-processing" strategies under pressure.

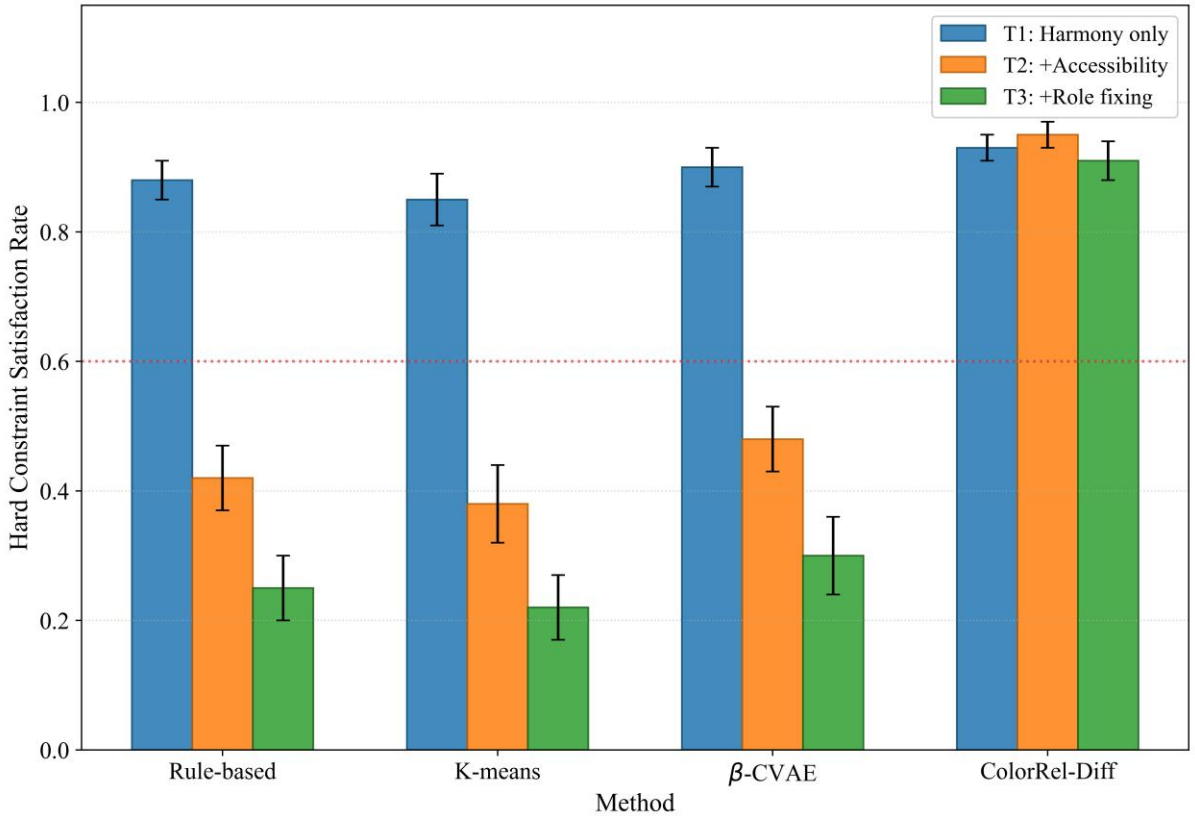



**Figure 2. Hard constraint satisfaction rates and confidence intervals of the bootstrap method for each method on the three increasing constraint strictness test sets**

The results in T3 shed further light on the value of the conflict softening mechanism. Under the joint constraints of role fixation and accessibility, the feasible region is the narrowest, and there are structural conflicts between some brand color and contrast requirements themselves. Under this condition, the direct failure strategy will return null results or a large number of invalid schemes, while the grading softening mechanism will output the minimum violation scheme with conflict hints. Figure 3 presents the output directions of each method on T3 in the form of Sanji: the left side is the input

scheme, the middle is the three directions of "fully satisfied", "output after softening" and "direct failure", and the right side is the distribution of constraint violation types of softening output. It can be seen that most of the schemes of the method in this paper enter into "full satisfaction" or "output after softening," and the proportion of direct failures is very low. It is important to emphasize that the solution with softened output has a higher usability proxy score than the null result of the direct failure strategy. The significance of this comparison is not how good the softened solution is per se, but that "having available solutions and informing conflicts" is superior to "no output" in the design process, which forces the designer to manually adjust from the beginning, while the former at least provides a modifiable starting point.

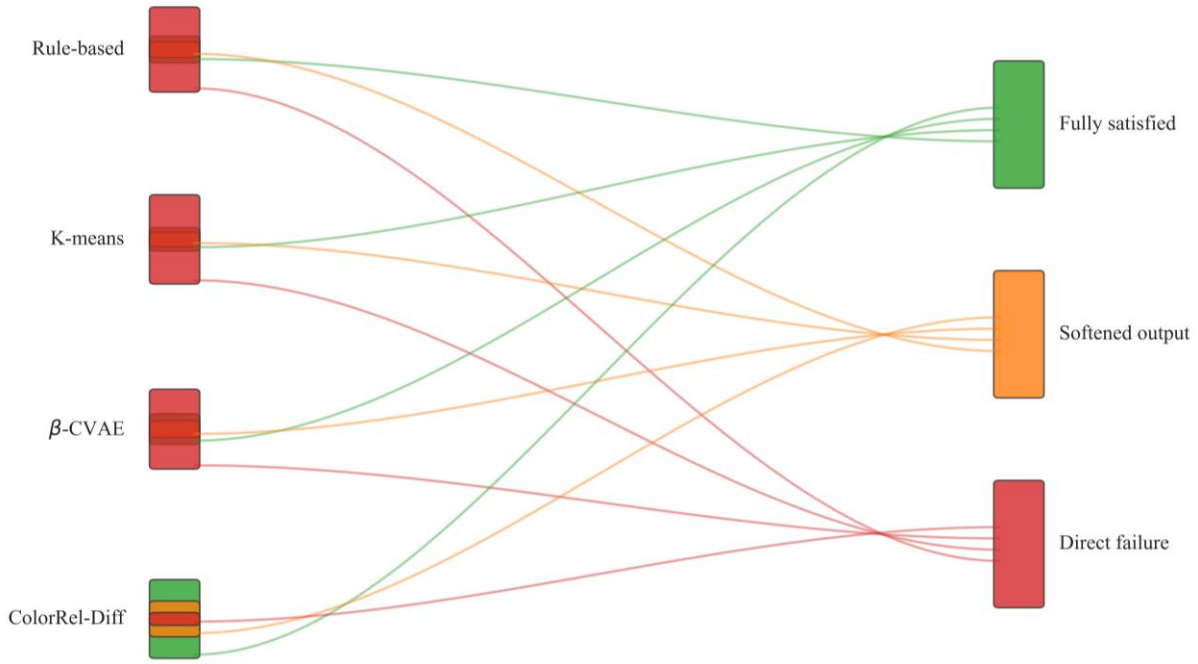



**Figure 3. Sankey decomposition of the output direction and softening violation type of each method under the joint constraint of T3**

### 5.3 Proxy evaluation: representative boundaries of automatic indicators

The task of layer 3 is not to prove that the proposed method wins in the ranking of experts, but to define the extent to which the automatic indicators of layer 1 and layer 2 can represent design judgment. Figure 4 presents the Spearman rank correlation coefficient between each automatic indicator and Bradley-Terry proxy score in the form of thermal map, and the rows and columns are the automatic indicator and the three evaluation dimensions respectively. The results show that the harmony score is significantly positively correlated with the dimension of "overall preference" and moderately correlated with the dimension of "visual hierarchy clarity", indicating that the automatic scoring based on harmony rules does capture a considerable part of the design judgment. The satisfaction rate of hard constraints is significantly correlated with the dimension of "scene suitability", which is consistent with the setting that accessibility constraints directly correspond to the functional requirements of the scene. However, the correlation between the diversity indicator and the three dimensions is weak and insignificant, which deserves attention: diversity is seen as a positive indicator in layer 1, but the proxy evaluation shows that it does not directly translate into design preferences, and excessive pursuit of diversity may produce "strange" combinations that are not available in design.

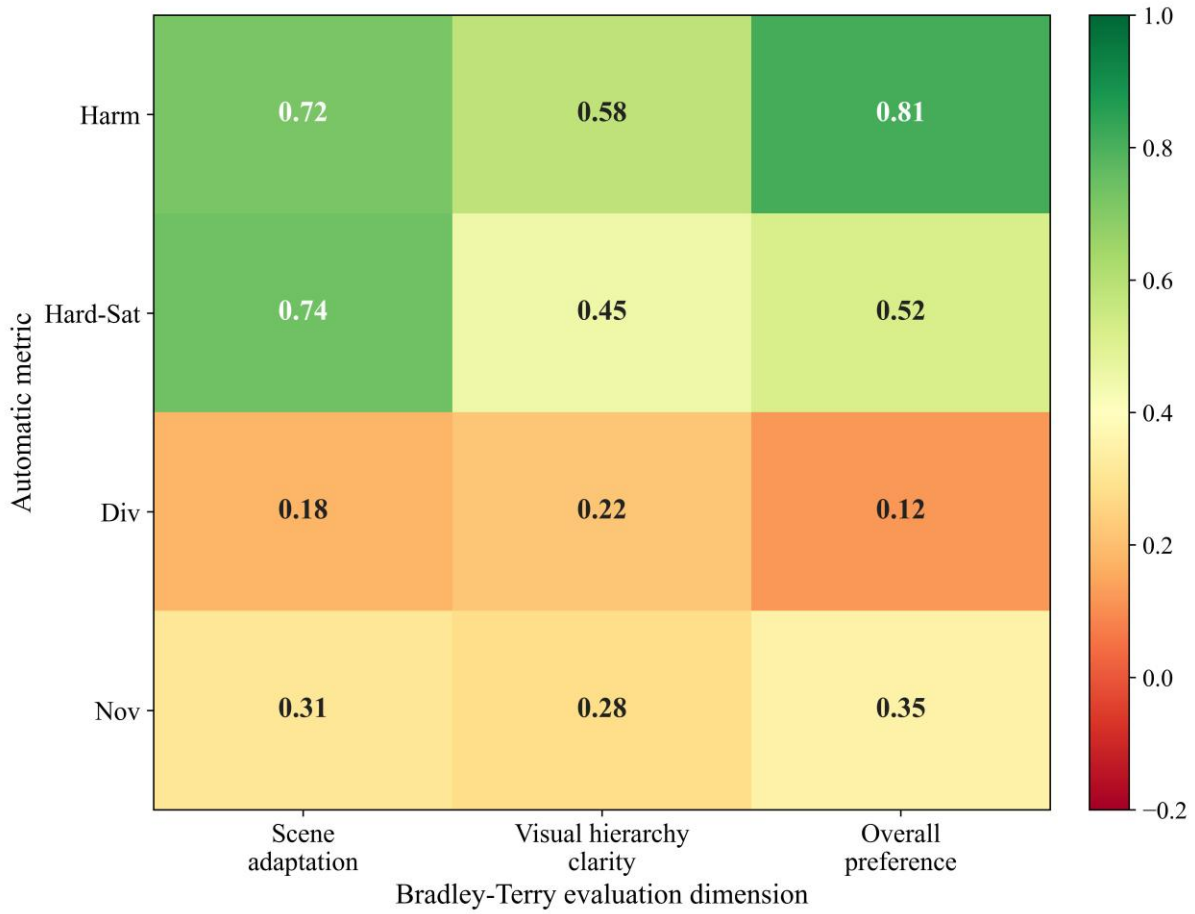



**Figure 4. Spearman rank correlation thermogram of automatic indicators and Bradley-Terry proxy scores**

The reason why harmony degree and constraint satisfaction rate can represent design judgment is that they directly correspond to the check items that designers actually perform when selecting colors, whether colors are coordinated and whether words are readable. Diversity describes the statistical coverage of a scheme set. When evaluating a single scheme, designers do not take "whether this scheme makes the set more diverse" as the judgment basis, and there is a granularity mismatch in the evaluation objects between the two. This mismatch suggests that if diversity is taken as the optimization goal, it needs to be positioned as the constraint at the set level rather than the goal at the scheme level, otherwise there will be a directional conflict with the design availability. Another finding of the proxy evaluation is that the softening output scheme scores no lower than the fully satisfied scheme in "scenario fit," and the difference is not significant within the confidence interval, indicating that constraint softening does not come at the cost of scenario fit, which is consistent with the design in which the accessibility bottom line is not shaken in the softening priority.

### 5.4 Ablation attribution: Asymmetric distribution of component contributions

The value of the ablation results is to decompose the overall performance into components and test whether the degradation pattern is consistent with the component function hypothesis. Figure 5 presents the retention rates of the four ablation configurations relative to the complete model in terms of diversity, harmony, T2 hard constraint satisfaction rate and T3 soft violation degree in the form of radar chart.

The axial outward direction of radar indicates that the higher the retention rate is and the less the degradation is. The degradation pattern shows an obvious asymmetric distribution: removing the relational coding module mainly leads to a decrease in the retention rate of harmony degree, while the retention rate of constraint satisfaction rate changes little, indicating that the function of this module focuses on the description of relationship quality rather than constraint compliance. Removing the constraint guidance term mainly leads to a decrease in the retention rate of soft constraint violation degree, while the hard constraint satisfaction rate is almost unaffected, which is completely consistent with the labor division hypothesis of "shape of soft guidance modulation distribution and hard mask truncation infeasible support set" in Section 2. The impact of removing course learning on layer 1 and layer 2 indicators is relatively mild, but the retention rate of T3 satisfaction rate decreases to a certain extent. Moreover, the training convergence curve shows that the loss plateau period occurs when the proportion of hard constrained samples increases in single-stage training, indicating that the effect of course learning is mainly reflected in the convergence stability and generalization under the scenario of small sample constraints, rather than the upper limit of the final distribution quality. Among the four ablation configurations, the complete model is in the outermost circle on all radar axes, and this joint mode rather than the single-axis lead is the evidence of the effective synergy of the components.

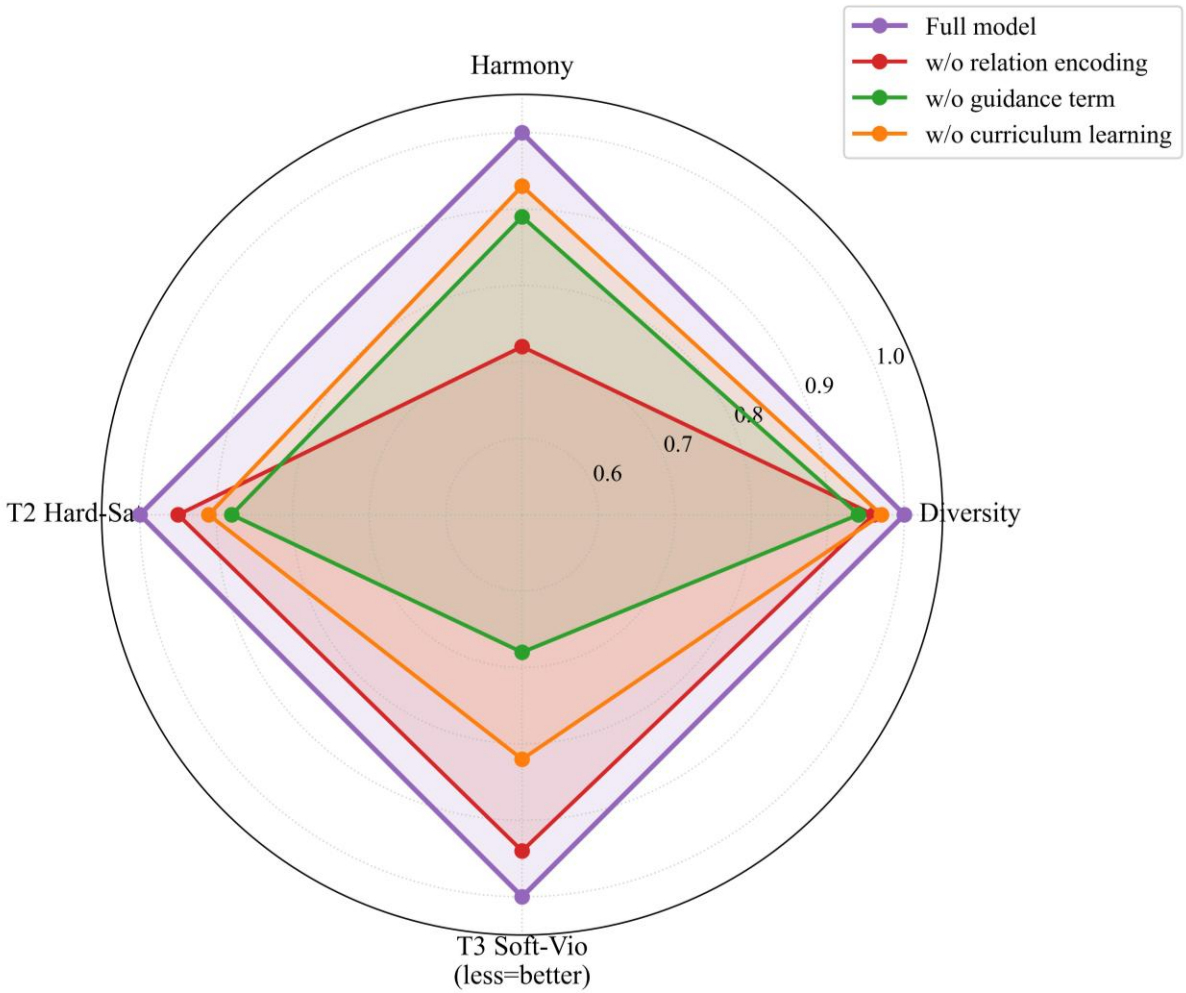



**Figure 5. Radar plot of index retention rates for the four ablation configurations relative to the complete model**

A more complete causal picture can be obtained by comparing the ablation results with the grading

mode of layer 2: the advantages of the proposed method in T2 and T3 mainly come from the coordination of the constraint guidance term and the hard mask, the former improving the compliance of soft constraints, and the latter ensuring that the hard constraints cannot be violated. The leading degree of harmony at the first level mainly comes from the relationship coding module. This means that if the application scenario is dominated by unconstrained concept exploration, the relational coding module is the core contribution point, and the marginal value of constrained guidance is limited. If the scenario is dominated by scheme generation with hard constraints, the constraint guidance and mask mechanism is the decisive component. This conditional dependence attribution conclusion is more engineering instructive than the blanket statement "all components matter" and answers the question raised at the beginning of this section: The advantage of ColorRel-Diff does not exist uniformly but is enlarged with the tightening of constraints. Its lead in the unconstrained interval is narrowed by strong baseline competition, and its real discrimination is reflected in the pressure interval where the feasible region is narrowed, and the post-processing strategy fails.

# 6. DISCUSSION AND DESIGN IMPLICATIONS

The results in Section 5 reveal a non-uniform dominance distribution: ColorRel-Diff's discrimination is concentrated in the pressure interval of constraint tightening, while its lead is limited in the unconstrained concept exploration interval. This distributional feature itself points to the more fundamental question of what role automatic color generation should play in the design process. If it is positioned as "alternative designer color selection," then the value of the method should be measured by whether the output scheme is directly adopted. A more reasonable positioning is that the algorithm provides a set of interpretable candidate schemes in the constraint space and shifts the designer's work from "starting from zero" to "judging and modifying in the candidate set". The algorithm is responsible for the constraint compliance check and feasible region exploration, and the designer is responsible for the final discretion of cultural semantics and brand narrative. This section discusses sources of interpretability around this positioning, the implications of failure modes, the division of processes with neighboring methods, and the boundaries of the current framework.

Explainability is a prerequisite for the establishment of this role positioning. If the algorithm outputs a set of colors that cannot be explained, the designer can neither trust nor effectively modify them, and the set of candidate solutions reduces to a random lottery. The interpretability of ColorRel-Diff is not an additional explanatory text but is built into the generation mechanism: The relationship consistency score calculates the relationship score pair by pair for the candidate colors at each step of sampling. This intermediate output can be traced back to "why this set of colors is recommended ", which pair of colors contributes a high score for the complementary relationship, which pair is truncated by the mask due to insufficient contrast, and which one cannot be adjusted due to fixed roles. Instead of seeing the final solution output from the black box, the designer sees a decision chain with a relationship basis, so that the modification has a direction: if the transition between two roles is not satisfied, the designer can target the specific relationship score item rather than the overall redraw. This feature is complementary to the conflict prompt mechanism in Section 3, where the latter tells "what is not satisfied" and the former tells "what is preferred", which together constitute the interface for designer intervention. It should be admitted that the current relationship score is based on formable reconciliation rules and contrast formulas, and its interpretability stops at the level of rules, which cannot explain why this set of colors fit the cultural temperament of a brand.

The analysis of failure modes further clarifies the boundaries of the role of algorithms. As observed

in Section 5, when the style reference palette is in serious conflict with the direction of functional semantic constraints, the generated results tend to sacrifice soft style constraints and preserve functional hard constraints. This prioritization behavior is a direct consequence of the grading softening strategy of Section 3, which is designed based on the readability and brand consistency related to the functional usability of the scheme, while the style orientation belongs to the adjustable aesthetic goal. However, from the perspective of design practice, it is worth careful discussion whether this priority is generally true: in some brand extension scenarios, style consistency itself is an unyielding constraint, and the "rational" retreat of the algorithm is contrary to the design intention. This shows that the priority setting in the grading softening strategy should not be regarded as a universal truth but should be given to the user as a configurable parameter. What the algorithm can do is to output the minimum violation scheme under the given priority and clearly inform the conflict, rather than deciding for the designer which kind of constraints are more important. Another implication of the failure mode is that the value of the algorithm in conflict scenarios is not to give the "right answer", but to make the conflict explicit. In manual color matching, designers often realize the structural contradiction between constraints after repeated attempts, but the algorithm can detect and prompt the conflict before sampling, which has more process value than the output of a single scheme.

The positioning difference with image coloring and palette transfer method needs to be understood from the stage division of the design process. The image coloring method serves to adjust the color matching the existing visual draft. The input is grayscale map or line draft, and the output is a colored image. The structure of the source map constitutes a strong constraint. The palette transfer method serves to transfer the style of the reference palette to the target image, and its core is the mapping of color distribution. Both of them are based on the premise of "existing visual carrier" and are in the middle and later part of the design process. The output of ColorRel-Diff is a role-based color scheme from scratch, and the input is a functional role, accessibility requirements and style reference. Colorrel-diff does not rely on any visual carrier and is in the early stage of the design process, concept conception and scheme exploration. This stage division of labor means that the three types of methods cannot be directly replaced: calling the coloring method in the scheme exploration stage lacks the input carrier, and calling the method in the visual draft adjustment stage loses the source graph structure information. Another implication of process division is the difference in evaluation criteria: The evaluation of the former method should focus on the coverage breadth and constraint compliance of the candidate set, while the evaluation of the latter method should focus on the coordination with the source graph structure. In Section 5, the layer three agent evaluation takes "scene suitability" and "overall preference" as the dimensions, which is designed for the former positioning. If the same evaluation system is directly applied to the coloring method, the conclusions are not comparable.

The boundary conditions of the current framework are first reflected in the degree of formalization of constraint types. The reason why functional semantics, accessibility and style orientation can be encoded is that they can be transformed into mathematical objects such as location prior, contrast threshold or set similarity. However, it is highly dependent on the color semantics of cultural context or brand narrative, such as the taboo meaning of a certain red in a specific culture and the historical inheritance of a brand color system, and there is still a lack of effective coding means. This kind of semantics cannot be reduced to hue Angle difference or contrast ratio, and it is difficult to fully express with a small number of reference color palates. At present, its processing can only rely on the manual screening of designers on the candidate set. The second layer of boundary is reflected in the number of roles in the generation unit, which is fixed as five functional roles in this paper. When the design scheme involves more detailed role division (such as multi-level text levels and coexistence of multiple

emphasis colors), the complexity of sequence length and role conditional modulation increases accordingly. Although the cost of offline calculation of word list scale and relationship diagram does not increase with the sequence length, However, the modeling ability of denoising network for long sequences has not been verified. The third layer of the boundary is reflected in the agency nature of the evaluation. The recognition of these boundaries does not diminish the value of the method in this paper in its application range, but makes the role of "algorithm provides interpretable candidate set, and the designer retains the final discretion" clearly applicable conditions: In the scheme exploration stage when constraints can be formalized, the number of roles is moderate, and the evaluation is dominated by directional judgment, the framework can undertake the functions of feasible region exploration and compliance inspection. Beyond this interval, the algorithm should return to the auxiliary position and return the judgment right to the design subject.

# 7. CONCLUSION

This paper proposes and verifies the ColorRel-Diff framework for automatic generation of color collocation in design schemes. The core of its methodology is not the introduction of diffusion model as a generation tool itself, but the promotion of the quality of the relationship between colors from an ex-post evaluation criterion to a first-class guiding signal in the generation process. The color relationship is always outside the generation target and only tested after the output due to the hard constraints of the generation path or the dependence on the harmonic template, or the conditions of the semantic label and the reference image. The work in this paper shows that when the relationship consistency score acts on the sampling distribution in the form of differentiable guide term and when the hard constraint truncates the infeasible support set in the form of feasibility mask, the generated results are significantly better than the post-processing filtering strategy in the pressure interval of constraint tightening. The grading pattern in Section 5 provides direct evidence for this judgment: advantages are not uniformly distributed across all conditions, but are amplified as the feasible region Narrows, which precisely shows that the value of relational guidance lies in embedding constraint compliance endogenously in the generation process, rather than passively screening it at the output. Matching with this architectural contribution is a three-layer test system, which describes different sides of the generated results by the distribution index of the set layer, the satisfaction rate of the constraint layer and the ranking score of the agent layer respectively, and defines the representative boundary of the automatic index by correlation analysis, so that the judgment of "good generated results" changes from the leading of a single index to the convergence of multi-level evidence.

The extensibility of the methodological contribution of this paper needs to be defined with clear premises. The effectiveness of ColorRel-Diff is based on two structural conditions: first, the generation object is discrete combination rather than continuous signal, so the discrete diffusion of mask and the token-level truncation of feasibility mask can be implemented naturally; Second, there is a formable relationship structure between the combination elements, so the relationship graph coding and consistency score gradient have a clear domain of definition. For discrete combination generation problems that meet these two conditions, the idea of "relational coding plus constraint guided denoising" can be used. The coordination relationship between word weight, literal and style in font collocation, the matching relationship between roughness, reflectivity and color in material combination, the relationship between area, hierarchy and contrast in layout element color matching, they can be encoded as edge features on the relationship graph and converted into guidance signals during the sampling period. On the contrary, when the relationship between the combination elements is highly dependent

on the context and difficult to formalize, the transfer value of this idea will be greatly weakened. This boundary condition is also an honest definition of the scope of this paper's contribution: this paper proposes a generative paradigm oriented to formable relational constraints, rather than universal color intelligence.

Future work can be extended in two directions. The first is the expansion of constraint types. The harmonious relationship, accessibility threshold and style similarity handled by the current framework are all manually defined rule constraints, and their expressive power is limited by the prior knowledge of the rule designer. Implicit aesthetic preference is incorporated into the guidance mechanism in a learnable form, such as learning the revision item of relationship score from the modification behavior of designers, so that the score head can continue to evolve with the use of data after the initialization of rules, which is expected to break through the upper limit of the expression ability of manual rules on cultural context and brand narrative. The second is the construction of interactive design loop. The method positioning of this paper is to provide an interpretable set of candidate solutions rather than replace the designer's decision. However, the current generation and evaluation is still a one-way process, and the designer's modification feedback has not yet flowed back to the model. Exploring the closed-loop mechanism of generation, evaluation, modification and regeneration, so that the intermediate output of relationship scoring can not only be used as the basis for interpretation, but also as the signal of preference learning, will be a key step in the framework from offline generation tool to design collaborative system. These two paths together point to a deeper problem: the ultimate goal of automatic color generation is not to eliminate the designer's color selection work, but to redistribute the human-machine division of labor between constraint compliance and meaning discretion. The work in this paper provides a testable technical starting point for this division of labor.

**Abbreviations**

AdaLN, Adaptive Layer Normalization;

CI, Confidence Interval;

CIELAB, CIE L*a*b* color space;

CVAE, Conditional Variational Autoencoder;

β-CVAE, β-Conditional Variational Autoencoder;

JSD, Jensen-Shannon Divergence;

LeakyReLU, Leaky Rectified Linear Unit;

WCAG, Web Content Accessibility Guidelines;

Div, Diversity;

Nov, Novelty;

w/o, without;

**Supplementary Material**

Not applicable.

## Appendix

Not applicable.

## Ethics approval and consent to participate.

This study did not involve human participants, animal subjects, or any data requiring ethical approval. Therefore, ethics approval and consent to participate are not applicable.


## Acknowledgements

The authors would like to thank the editors of this journal and all the anonymous reviewers who provided valuable comments on this work.


## Competing interests

The authors declare that they have no financial or personal relationships that may have inappropriately influenced them in writing this article.

## Author contributions

All authors have read and agreed to the published version of the manuscript. The author's contributions are specified as follows: **M.L.:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources. **Q.H.:** Conceptualization, Data Curation, Writing – original draft, Writing – review & editing, Visualization, Supervision, Project administration.


## Funding information

The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.


## Data availability

The data that support the findings of this study are available upon request from the corresponding authors, **Q.H.**

## Disclaimer

The views and opinions expressed in this article are those of the authors and are the product of professional research. It does not necessarily reflect the official policy or position of any affiliated

institution, funder, agency, or that of the publisher. The authors are responsible for this article's results, findings, and content.

**Declaration of AI and AI-assisted Technologies in the Writing Process**

During the writing of this article, the author used ChatGPT for spelling and grammar checking. After using this tool, the author reviewed and edited the content as needed and assumes full responsibility for the final published content.

## REFERENCES


[1] Feng, L., Yu, G., Miao, M., & Sun, J. (2025). Sustainable Color Development Strategies for Ancient Chinese Historical Commercial Areas: A Case Study of Suzhou's Xueshi Street–Wuzounfang Street. *Sustainability*, *17*(11), 4756. https://doi.org/10.3390/su17114756

[2] Kim, B., Lee, G., & Suk, H. J. (2026). Color adjustment of brand logos for dark mode display. *Plos one*, *21*(1), e0339392. https://doi.org/10.1371/journal.pone.0339392

[3] Zhou, X., Xiao, C., Yoon, S., & Zhu, H. (2026). The color of status: color saturation, brand heritage, and perceived status of luxury brands. *Journal of Consumer Research*, *52*(6), 1232-1252. https://doi.org/10.1093/jcr/ucaf029

[4] Zhang, H., Song, H., Li, S., Zhou, M., & Song, D. (2023). A survey of controllable text generation using transformer-based pre-trained language models. *ACM Computing Surveys*, *56*(3), 1-37. https://doi.org/10.1145/3617680

[5] Luo, J., Tian, Y., & Wang, Z. (2024). Research on unmanned aerial vehicle path planning. *Drones*, *8*(2), 51. https://doi.org/10.3390/drones8020051

[6] Museros, L., Sanz, I., Falomir, Z., & Gonzalez-Abril, L. (2020). Creating, interpreting and rating harmonic colour palettes using a cognitively inspired model. *Cognitive Computation*, *12*(2), 442-459. https://doi.org/10.1007/s12559-018-9589-2

[7] Zhang, Y., Lin, C., Tang, S., Chen, H., Zhou, S., Ma, Y., & Tresp, V. (2025, November). Swarmagentic: Towards fully automated agentic system generation via swarm intelligence. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing* (pp. 1778-1818). https://doi.org/10.18653/v1/2025.emnlp-main.93

[8] Liu, S., Tao, M., Huang, Y., Wang, C., & Li, C. (2022). Image-driven harmonious color palette generation for diverse information visualization. *IEEE Transactions on Visualization and Computer Graphics*, *30*(7), 3089-3103. https://doi.org/10.1109/tvcg.2022.3226218

[9] Gao, Y., Liang, J., & Yang, J. (2025). Color palette generation from digital images: A review. *Color Research & Application*, *50*(3), 250-265. https://doi.org/10.1002/col.22975

[10] Sun, M. H., & Chou, T. R. (2026, February). Disentangled Emotion-Controllable Color Palette Generation via β-CVAE with Conditional Prior. In *2026 International Conference on Artificial Intelligence in Information and Communication (ICAIIC)* (pp. 605-610). IEEE. https://doi.org/10.1109/icaiic68212.2026.11454419

[11] Wang, K., Jiang, Z., Feng, H., Zhao, W., Liu, L., Li, J., ... & Lin, W. (2026, July). CreditDecoding: Accelerating Parallel Decoding in Diffusion Large Language Models with Trace Credit. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)* (pp. 11105-11123). https://doi.org/10.18653/v1/2026.acl-long.509

[12] Araujo Oliveira, R. M., Oliveira Sant'Anna, Â. M., Ferreira, P. H., & Purcino de Souza, E. E. (2026). Fuzzy hybrid system based on ensemble learning for classification in intelligent manufacturing processes. *The International Journal of Advanced Manufacturing Technology*, 1-29. https://doi.org/10.1007/s00170-026-18928-3

[13] Hristov, H., Glushkova, T., Cheresharov, S., & Stoeva, M. (2022, October). A model for designing accessible color and contrast for users with visual deficiency and color blindness. In *2022 IEEE 11th international conference on intelligent systems (IS)* (pp. 1-7). IEEE. https://doi.org/10.1109/is57118.2022.10019637

[14] Lewandowska, A., Olejnik-Krugly, A., Jankowski, J., & Dziśko, M. (2021). Subjective and objective user behavior disparity: Towards balanced visual design and color adjustment. *Sensors*, *21*(24), 8502. https://doi.org/10.3390/s21248502

[15] Sun, C., Li, C., Lin, X., Zheng, T., Meng, F., Rui, X., & Wang, Z. (2023). Attention-based graph neural networks: a survey: C. Sun et al. *Artificial intelligence review*, *56*(Suppl 2), 2263-2310. https://doi.org/10.1007/s10462-023-10577-2

[16] Guo, X., Yu, Z., & Yin, J. (2026). Fusion of Radiomics and Gated Graph Attention Network for Pulmonary Nodule Malignancy Classification. *Journal of Imaging*, *12*(9), 408. https://doi.org/10.3390/jimaging12090408

[17] Höllerer, S., Papaxanthos, L., Gumpinger, A. C., Fischer, K., Beisel, C., Borgwardt, K., ... & Jeschek, M. (2020). Large-scale DNA-based phenotypic recording and deep learning enable highly accurate sequence-function mapping. *Nature communications*, *11*(1), 3551. https://doi.org/10.1038/s41467-020-17222-4

[18] Elfmann, C., & Stülke, J. (2023). PAE viewer: a webserver for the interactive visualization of the predicted aligned error for multimer structure predictions and crosslinks. *Nucleic acids research*, *51*(W1), W404-W410. https://doi.org/10.1093/nar/gkad350

[19] Meng, C., Gao, J., Tian, Y., Chen, H., Zhang, W., & Luo, R. (2024). Stochastic solutions for simultaneous seismic data denoising and reconstruction via score-based generative models. *IEEE Transactions on Geoscience and Remote Sensing*, *62*, 1-15. https://doi.org/10.1109/tgrs.2024.3421597

[20] Garber, T., & Tirer, T. (2024, June). Image restoration by denoising diffusion models with iteratively preconditioned guidance. In *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (pp. 25245-25254). IEEE. https://doi.org/10.1109/cvpr52733.2024.02385

[21] Valenko, S. I., Čačić, M., Stanimirović, I. Ž., & Zorko, A. (2025). Visual Perception of Peripheral Screen Elements: The Impact of Text and Background Colors. *Applied Sciences*, *15*(14), 7636. https://doi.org/10.3390/app15147636

[22] Jakovljević, T., Janković, M. M., Savić, A. M., Soldatović, I., Todorović, P., Jere Jakulin, T., ... & Ković, V. (2021). The sensor hub for detecting the developmental characteristics in reading in children on a white vs. colored background/colored overlays. *Sensors*, *21*(2), 406. https://doi.org/10.3390/s21020406

[23] Collins, M., Vice, J., French, T., & Mian, A. (2026, April). Natadiff: Adversarial boundary guidance for natural adversarial diffusion. In *International Conference on Learning Representations* (Vol. 2026, pp. 12099-12145). https://doi.org/10.48550/arXiv.2505.20934

[24] Kabir, A. M., Thakar, S., Malhan, R. K., Shembekar, A. V., Shah, B. C., & Gupta, S. K. (2021). Generation of synchronized configuration space trajectories with workspace path constraints for an ensemble of robots. *The International Journal of Robotics Research*, *40*(2-3), 651-678. https://doi.org/10.1177/0278364920988087

[25] Wainer, J. (2023). A Bayesian Bradley-Terry model to compare multiple ML algorithms on multiple data sets. *Journal of Machine Learning Research*, *24*(341), 1-34. https://doi.org/10.48550/arXiv.2208.04935

[26] Fayomi, A., Majeed, R., Algarni, A., Akhtar, S., Jamal, F., & Nasir, J. A. (2022). Forecasting tennis match results using the Bradley-Terry model. *International Journal of Photoenergy*, *2022*(1), 1898132. https://doi.org/10.1155/2022/1898132

[27] Fageot, J., Farhadkhani, S., Hoang, L. N., & Villemaud, O. (2024, March). Generalized bradley-terry models for score estimation from paired comparisons. In *Proceedings of the AAAI Conference on Artificial Intelligence* (Vol. 38, No. 18, pp. 20379-20386). https://doi.org/10.1609/aaai.v38i18.30020

## APA Citation

Liu, M., & He, Q. (2026). Relation-Aware and Constraint-Guided Discrete Diffusion for Design Color Scheme Generation. *Journal of Discovery Core, 1*(2), 80-109. **https://doi.org/10.67541/jdc2610**