

When Rubrics Fail: Error Enumeration as Reward in Reference-Free RL Post-Training for Virtual Try-On

Wisdom Ikezogwo^{1 2} Mehmet Saygin Seyfioglu² Ranjay Krishna^{1 3} Karim Bouyarmane²




Figure 1. **Qualitative comparison of reward designs on non-flat references:** Given the input, IEC avoids *attribute preservation* errors across 3 examples, *garment transfer* failures in example C, *source integrity* issues in A and D, and *realism* errors in examples B and C, over all other reward designs. Per-category radar plots are in Appendix B.3.

Abstract

Reinforcement learning with verifiable rewards (RLVR) and Rubrics as Rewards (RaR) have driven strong gains in domains with clear correctness signals and even in subjective domains by synthesizing evaluation criteria from ideal reference answers. But many real-world tasks admit multiple valid outputs and lack the single ideal answer that rubric generation depends on. We identify this *reference-free* setting as a gap in current post-training methods and propose **Implicit Error Counting (IEC)** to fill it. Instead of checking what a response gets right against a rubric, IEC enumerates what it gets wrong, applying severity-weighted scores across task-relevant axes and converting them into calibrated per-aspect rewards. We show that naïve explicit enumeration is too

noisy for stable optimization, and that two design choices: implicit score emission and group calibration are necessary to make error counting a reliable reward. As a case study, we validate IEC on virtual try-on (VTO), a domain that is simultaneously too constrained for holistic scoring and too permissive for rubric-based evaluation: subtle garment errors are unacceptable, yet many output variations are correct. We introduce Cascaded Error Counting (CEC) as an evaluation metric, which tracks human preferences well (60% top-1 vs. 30% others), and curate Mismatch-DressCode (MDressBench), a benchmark with maximal attribute mismatch to stress-test reward designs. On MDressBench, IEC outperforms RaR across all metrics (CEC: 5.31 vs. 5.60 on flat references; 5.20 vs. 5.53 on non-flat). On VITON-HD and DressCode, IEC matches or surpasses six baselines on 6 of 8 perceptual metrics. These results suggest that when ideal answers are unavailable, counting errors provide a stronger signal than constructing rubrics.

¹University of Washington, Seattle, USA ²Amazon, Seattle, USA ³Allen Institute for AI. Correspondence to: Mehmet Saygin Seyfioglu <mseyfiog@amazon.com>.

1. Introduction

Reinforcement learning has become a standard post-training step for aligning generative models with human expectations. In domains with verifiable outcomes—mathematics, code generation, game playing—reward is straightforward: a scoring function or test environment confirms correctness (Shao et al., 2024; Ouyang et al., 2022). Reinforcement Learning with Verifiable Rewards (RLVR) (Shao et al., 2024; Rafailov et al., 2023; Wallace et al., 2024; Black et al., 2023; Fan et al., 2024; Xue et al., 2025) produces dramatic gains precisely because the reward signal is cheap, unambiguous, and dense.

Real-world tasks are rarely so clean. Medical diagnosis, scientific reasoning, and creative generation require nuanced, multi-criteria evaluation with no single correct answer. Rubrics as Rewards (RaR) (Gunjal et al., 2025) addresses this gap: given a prompt and an *ideal reference answer*, it synthesizes instance-specific rubric criteria and scores responses against them. RaR provides structured, interpretable supervision that outperforms coarse Likert scoring. Its core assumption, however, is that an ideal answer exists and can ground rubric generation. When it does, as in medical QA with expert-written references, rubrics capture subtle quality dimensions that binary correctness cannot.

But many domains lack ideal answers altogether. In embodied control, multiple action sequences can achieve the same goal. In open-ended design tasks, “correctness” is defined not by what the output matches, but by what it avoids—wrong colors, broken geometry, identity drift. We refer to this as the **reference-free** setting: the task is subjective enough that rubrics are needed, yet no ideal output exists to generate them from. In this setting, the nature of quality assessment inverts. Rather than asking “does this response satisfy a checklist derived from an ideal answer?” one asks “what specific errors does this response exhibit?” When multiple outputs are plausible, enumerating failures is more stable than enumerating successes, because the space of errors is typically smaller and more structured than the space of valid completions.

We test this hypothesis with **Implicit Error Counting (IEC)**, a reward framework for reinforcement post-training in reference-free domains. Given a candidate output and its conditioning inputs, a multimodal judge enumerates errors across task-relevant evaluation axes, weights them by severity, and emits calibrated per-aspect scores. The “implicit” qualifier is critical: asking the judge to expose raw error lists (Explicit Error Counting, or EEC) introduces high variance from inconsistent enumeration, where near-identical outputs receive different error counts due to surface-form variation in the judge’s language. IEC instead asks the judge to *internalize* the counting and output only scores, preserving the conceptual grounding of error enumeration while stabi-

lizing the reward signal. A lightweight group calibration step further reduces prompt-to-prompt scale drift. Notably, IEC is also more efficient: it requires a single judge call per candidate, whereas RaR requires two (rubric generation followed by evaluation).

We validate IEC on virtual try-on (VTO) as a case study that exemplifies the reference-free challenge. VTO synthesizes images of users wearing target garments (Choi et al., 2021; Morelli et al., 2022; Zhu et al., 2023; Li et al., 2025b; Seyfioglu et al., 2024), and its failure modes are simultaneously subtle and consequential: a slightly wrong sleeve length, a discontinuous pattern, or a shifted neckline can render a result unusable. Crucially, many distinct outputs are correct; the garment can drape differently, lighting can vary, and fine details can differ. So no single ideal answer grounds rubric generation. At the same time, errors are *enumerable* and *localizable*: they fall along well-defined axes (garment transfer, attribute preservation, realism, lighting, source integrity) and can be detected by multimodal judges. Supervised objectives on *narrow* datasets (Morelli et al., 2022; Choi et al., 2021) cannot cover the combinatorial space of poses, body shapes, textures, and garment types, and tend toward shortcut learning (Labs et al., 2025; Wu et al., 2025; UniWorld Team, 2025), especially copy-paste behavior in masked-image editing. This motivates on-policy post-training that improves behavior without new paired supervision.

To stress-test reward designs under distribution shift, we curate **Mismatch-DressCode (MDressBench)**, a benchmark of 700 source-reference pairs selected for maximal attribute disagreement (e.g., short-sleeve source with long-sleeve reference). We also introduce **Cascaded Error Counting (CEC)**, an evaluation metric that builds a shared error vocabulary across candidates, enabling consistent measurement without ground truth. In a human study validating CEC as a selection metric, it achieves 60% top-1 accuracy in recovering annotator preferences, compared to 30% for both direct scoring and RaR-based selection.

Across MDressBench, IEC consistently outperforms both direct scoring and RaR on all metrics, achieving relative improvements of 5.96% (CEC), 2.32% (garment transfer), 3.48% (attribute preservation), 4.33% (realism), 3.3% (lighting), and 0.6% (source integrity) over RaR on non-flat references. On VITON-HD and DressCode, IEC outperforms all supervised and RL baselines on 6 of 8 perceptual metrics. Our ablations confirm that EEC underperforms IEC, and that group calibration provides consistent gains.

2. From Rubrics to Error Counting

We first formalize the gap between rubric-based and reference-free reward settings (§2.1), then present error

counting as a general reward framework (§2.2), and finally instantiate it for virtual try-on (§2.3).

2.1. The Reference-Free Gap

Rubrics as Rewards (Gunjal et al., 2025) defines a structured reward function for each prompt x via a set of k rubric items $\{(w_j, c_j)\}_{j=1}^k$, where w_j is the weight of criterion j and $c_j : (x, \hat{y}) \mapsto \{0, 1\}$ is a binary check. The reward is:

$$r_{\text{rubric}}(x, \hat{y}) = \frac{\sum_{j=1}^k w_j \cdot c_j(x, \hat{y})}{\sum_{j=1}^k w_j} \quad (1)$$

Rubric generation requires a reference answer y^* to ground the criteria: each c_j encodes a property that a high-quality response should satisfy, derived from analyzing y^* . When y^* exists (e.g., expert-written medical answers), this yields precise, instance-specific supervision.

The reference-free setting arises when three conditions hold simultaneously:

1. **Multiple valid outputs exist.** Many responses $\hat{y}$ may be equally correct, so no single y^* captures what “good” means.
2. **Quality is defined by the absence of errors.** An output is good if it follows the conditioning and has no observable errors, not because it matches a particular reference.
3. **Errors are enumerable and localizable.** Failure modes cluster along well-defined axes and can be detected by inspection.

These conditions are satisfied whenever the space of acceptable outputs is large but the space of errors is structured, a property shared by visual generation, embodied control, and open-ended design tasks. Under these conditions, rubric generation from a reference answer either produces criteria that are too generic (because no single y^* captures the range of valid outputs) or too specific (because y^* encodes one valid solution and penalizes equally valid alternatives). Error counting exploits the structure of these domains by inverting the evaluation: instead of scoring how much an output resembles an ideal, it scores how few detectable errors it contains.

2.2. Error Counting as a Reward Framework

Let x denote the conditioning input and $\hat{y}$ a candidate output. We define A task-relevant evaluation axes $\{a_1, \dots, a_A\}$. For each axis a , a judge function $f_a(x, \hat{y})$ returns a score in $[0, 1]$ by enumerating errors along that axis and mapping error counts to a monotonically decreasing score.

Explicit Error Counting (EEC). The judge produces a list of errors $\mathcal{E}_a = \{(e_1, s_1), \dots\}$ for each axis, where e_i is an error label and $s_i \in \{\text{minor}, \text{major}\}$ is its severity. The axis score is computed as a deterministic function of the weighted count:

$$R_a^{\text{explicit}} = g\left(\sum_{(e,s) \in \mathcal{E}_a} w(s)\right) \quad (2)$$

where g is a monotonically decreasing mapping and $w(\cdot)$ assigns severity weights. This is conceptually clean but suffers from high variance: non-deterministic judges produce different error labels for near-identical outputs, destabilizing within-group rewards. Appendix B.6 illustrates this with a concrete training group where EEC assigns a perfect score to a generation with visible errors while penalizing better candidates. See Appendix D.3 for prompt details.

Implicit Error Counting (IEC). We instead instruct the judge to *internally* count errors for each axis and output only a calibrated score $R_a^{(i)} \in [0, 1]$, along with a short error summary for interpretability. Each error monotonically reduces the axis score; catastrophic errors exponentially reduce it toward zero. We define *catastrophic* errors as wrong garment type, identity change, or a missing major garment region (see Appendix B.5 / D.1). The judge implicitly determines the per-error penalty weight. The aggregate reward is $R_i^{\text{implicit}} = \frac{1}{|A|} \sum_{a=1}^{|A|} R_a^{(i)}$.

The key distinction from EEC is that the reward depends on the judge’s *internal* assessment rather than on parsed error labels, making it robust to surface-form variation of the judge, and brittleness of fixed weights (Gunjal et al., 2025). The error summary provides interpretability without coupling directly to reward computation.

Group Calibration. Reward magnitudes vary across prompts due to difficulty and evaluator scale, introducing unnecessary variance in RL updates. We apply group-wise calibration using robust statistics. For a group reward vector $\mathbf{R}$:

$$z_i = \frac{R_i - \text{median}(\mathbf{R})}{1.4826 \cdot \text{MAD}(\mathbf{R}) + \varepsilon} \quad (3)$$

$$R_i^{\text{cal}} = \text{floor} + (1 - \text{floor}) \cdot \sigma(\alpha \cdot z_i) \quad (4)$$

where MAD is the median absolute deviation, 1.4826 rescales to standard-deviation scale under Gaussian assumptions, floor ≈ 0.05 prevents zero rewards, and σ is the logistic function with steepness $\alpha = 1$. This monotone transform preserves within-group ordering while reducing prompt-to-prompt scale variability.

2.3. Instantiation: Virtual Try-On

We formulate VTO as an image-conditioned inpainting problem. Given a source image I_s , reference/target garment I_r , mask I_m , and pose I_p , the model generates I_t where I_r appears on I_s while preserving identity, pose, and background.

Architecture. We build on a rectified-flow backbone (Liu et al., 2022; Lipman et al., 2022), following the token concatenation scheme of DiT-VTON (Li et al., 2025b). Each image is encoded into a latent space, patchified, and concatenated: $c = x_t \circ x_r \circ x_e$, where x_t is the noisy latent, x_r is the reference latent, and x_e is the masked source with pose stitched on. The flow-matching objective trains velocity $v_\theta(x_t, t, c)$ to approximate $v = x_1 - x_0$.

GRPO Post-Training. We apply Group Relative Policy Optimization (Shao et al., 2024) adapted to flow models (Xue et al., 2025; Liu et al., 2025). For each condition c , we sample K candidates and compute group-normalized advantages:

$$\hat{A}_i = \frac{R(x_i, c) - \text{mean}(\{R(x_j, c)\}_{j=1}^K)}{\text{std}(\{R(x_j, c)\}_{j=1}^K)} \quad (5)$$

Evaluation Axes. For VTO, we define $A = 5$ axes: (i) *garment transfer* (placement, sleeve/hem length, neckline, proportions); (ii) *attribute preservation* (color, pattern, texture); (iii) *realism* (drape, boundaries, halos, distortions); (iv) *lighting consistency*; and (v) *source integrity* (face, hair, background unchanged). Each axis is scored by the IEC judge (Equation 2.2), and the aggregate reward drives GRPO updates.

2.4. Baseline Rewards

We compare IEC against two standard reward designs, both evaluated by the same multimodal judge.

Direct Scoring. The judge outputs a scalar $R_i^{\text{direct}} \in [0, 1]$ for overall quality, guided by scoring bands: $[0.8, 1.0]$ for strong results, $[0.5, 0.8)$ for mixed quality, $[0.2, 0.5)$ for poor results, $[0.0, 0.2)$ for failures. It is efficient; however, it loses a fine-grained sensitivity to perceivable errors. See Appendix D.1.

Rubrics-as-Rewards (RaR). Following (Gunjal et al., 2025), we generate a prompt-specific rubric with criteria across categories {Essential, Important, Optional, Pitfall} with weights w_j . The judge evaluates each criterion with binary pass/fail. Let $\mathcal{P}_i = \{j : p_{ij} = 1, w_j > 0\}$ denote satisfied items and $\mathcal{F}_i = \{j : p_{ij} = 0, w_j < 0\}$ violated

pitfalls:

$$R_i^{\text{rubric}} = \text{clip}\left(\frac{\sum_{j \in \mathcal{P}_i} w_j + \sum_{j \in \mathcal{F}_i} w_j}{\sum_{j: w_j > 0} w_j}; 0, 1\right) \quad (6)$$

This requires two judge calls (rubric generation + evaluation) and, crucially, requires generating rubrics *without an ideal reference answer* in the VTO setting.

Reward Cost. IEC and Direct scoring each require a single judge call per candidate, whereas RaR requires two (rubric generation followed by evaluation). For a group of $K=12$ candidates, this means IEC uses 12 judge calls per training step compared to 24 for RaR, a $2\times$ reduction in judge compute, while achieving stronger performance across all metrics.

2.5. Cascaded Error Counting as a Metric

Plain error counting exhibits high variance as a metric because the label space drifts across images. Cascaded Error Counting reduces this by sharing an evolving error vocabulary across candidates within a group. Without pooling, a judge may enumerate zero errors for one candidate and seventeen for a near-identical one (see Appendix B.6); re-evaluating against a shared pool forces the judge to reconsider missed errors. We process candidates in blocks:

1. **Pool Phase:** The judge returns preliminary per-image error sets $\mathcal{E}_{i \in B} = \{(e_i, \text{sev}), \dots\}$. We canonicalize and merge into a pool $\mathcal{D}^{(0)}$ with semantic duplication.
2. **Verification Phase:** We re-evaluate images conditioned on $\mathcal{D}^{(0)}$, instructing the judge to find applicable pool errors and verify prior detections. This yields refined error sets $\mathcal{E}_i^{(1)}$.

Error canonicalization applies deterministic normalization: lowercasing, removing non-alphanumeric characters, and synonym replacement. A cyclical refinement pass further stabilizes the vocabulary. Final scores are weighted counts $C_i = \sum_{(e, \text{sev}) \in \mathcal{E}_i} w(\text{sev})$ with $w(\text{major}) > w(\text{minor})$. See Appendix D.4 for details.

3. Experiments

We evaluate error counting as a reward signal using virtual try-on as our primary testbed. Below, we describe the benchmark, baselines, and evaluation protocol.

Mismatch-DressCode (MDressBench): Existing VTO benchmarks randomly sample source-reference pairs within categories, yielding many pairs with similar attributes that simplify the transfer task (Morelli et al., 2022). To expose the failure modes and out-of-distribution artifacts typically




Figure 2. Training dynamics for four reward designs during GRPO post-training. **Left 2×2:** Per-method curves showing raw rewards (faded) and EMA (solid, $\alpha = 0.99$). **Right:** Normalized EMA showing relative improvement. IEC achieves smooth, monotonic improvement. Direct and RaR show high variance. EEC regresses early, confirming that explicit enumeration destabilizes optimization.

seen during model deployment, we curate MDressBench (Figure 3) from DressCode (train set) by selecting pairs with *maximal attribute disagreement*: short-sleeve sources with long-sleeve references, textured short dresses with ankle-length smooth references, and so on. We construct 10K pairs for RL training and reserve 350 pairs with the largest mismatch for evaluation, split across three categories (Top: 79, Lower: 122, Dress: 149) and two reference types (Flat and Non-flat), yielding 700 benchmark samples. For each pair, we follow DiT-VTON (Li et al., 2025b) to extract pose and inpainting masks for the source image, and align the reference garment crop. See Appendix A for details.

Baseline Finetuning

Dataset We train the baseline VTO model on a data mixture of VITON-HD (Han et al., 2018; Choi et al., 2021) and DeepFashion¹, covering a range of poses, body shapes, viewpoints, and garment styles.

Implementation Details

Our model builds on Kontext-dev (Labs et al., 2025), a 12B rectified-flow backbone, fine-tuned for 6.2k steps on the baseline mixture using the DiT-VTON token concatenation scheme (Li et al., 2025b). Training uses resolution 512×512 , batch size 8, learning rate 5×10^{-5} , AdamW, gradient checkpointing, and bf16 precision on 8 H100 (80GB) GPUs. Complete hyperparameter settings are listed



Figure 3. **MDressBench Examples.** Flat references (left) show the garment as a lay-flat product image on a neutral background. Non-flat references (right) show the garment worn by a model, requiring the model to “extract” the garment from arbitrary poses and backgrounds.

in Appendix B.1.

RL Post-Training. We apply GRPO with $K = 12$ generations per group and best-of- n down-selection ($n = 8$) to increase reward variance. We optimize for 60 steps with batch size 3, gradient accumulation 4, and bf16 precision. Sampling uses 30 steps at 512×512 with guidance scale 2.5, so we capture fine-grained errors that show up in the latter stages of sampling, e.g., color/pattern mismatch; we use 50% of steps for optimization biased toward earlier sampling steps. All reward computations use GPT-5-mini. Each reward design uses one judge call per candidate, except RaR (two calls: rubric generation + evaluation). To avoid self-referential bias, we use GPT-5 for evaluations.

Evaluation Protocol For each prompt, we generate 8 outputs per model and average metrics over two runs for stability, making the effective benchmark size 5600 samples. We evaluate robustly with multiple metrics, including: (i) CEC ($\downarrow$) as our primary metric (§2.5); (ii) RaR score ($\uparrow$); (iii) Direct score ($\uparrow$); and (iv) its aspect scores ($\uparrow$) for Garment Transfer, Attribute Preservation, Realism, Lighting, and Source Integrity. Additionally, we evaluate on the standard paired test splits of VITON-HD and DressCode using conventional perceptual metrics (LPIPS, SSIM, FID, KID) to assess generalization beyond MDressBench (Table 6).

Human Evaluation of CEC. To validate that CEC reflects human preferences as a *selection metric*, we conduct a 25-sample study where annotators rank try-on outputs from the baseline model across 8 images generated per sample, then evaluate each group using Direct, RaR, and CEC scoring and report top- k ranking accuracy. This study validates CEC as a metric, not IEC as a training reward; a dedicated human preference study of IEC outputs at scale is left to future work. Prior work on decomposed evaluation supports the principle that counting structured violations correlates

¹<https://www.kaggle.com/code/rkuo2000/deepfashion-tryon>

Table 1. Quantitative comparison on MDressBench (**flat** references). All metrics averaged across garment categories. $\uparrow$: higher is better; $\downarrow$: lower is better. Best values in **bold**. IEC outperforms all baselines across all eight metrics.

Method	CEC $\downarrow$	RaR $\uparrow$	Direct $\uparrow$	Garm. $\uparrow$	Attr. $\uparrow$	Real. $\uparrow$	Light. $\uparrow$	Src.Int. $\uparrow$
Baseline	6.018 $\pm$ 1.73	0.760 $\pm$ 0.14	0.796 $\pm$ 0.06	0.793 $\pm$ 0.16	0.815 $\pm$ 0.13	0.735 $\pm$ 0.12	0.760 $\pm$ 0.10	0.914 $\pm$ 0.08
Direct	5.373 $\pm$ 1.57	0.854 $\pm$ 0.11	0.840 $\pm$ 0.04	0.864 $\pm$ 0.09	0.873 $\pm$ 0.08	0.785 $\pm$ 0.08	0.771 $\pm$ 0.09	0.925 $\pm$ 0.07
RaR (Gunjal et al., 2025)	5.596 $\pm$ 1.69	0.826 $\pm$ 0.14	0.833 $\pm$ 0.05	0.858 $\pm$ 0.09	0.858 $\pm$ 0.09	0.771 $\pm$ 0.09	0.764 $\pm$ 0.10	0.929 $\pm$ 0.06
IEC (Ours)	5.312$\pm$1.72	0.875$\pm$0.09	0.854$\pm$0.04	0.873$\pm$0.08	0.886$\pm$0.07	0.803$\pm$0.07	0.794$\pm$0.07	0.933$\pm$0.06

Table 2. Quantitative comparison on MDressBench (**non-flat** references). All metrics averaged across garment categories. $\uparrow$: higher is better; $\downarrow$: lower is better. Best values in **bold**. IEC outperforms all baselines across all eight metrics.

Method	CEC $\downarrow$	RaR $\uparrow$	Direct $\uparrow$	Garm. $\uparrow$	Attr. $\uparrow$	Real. $\uparrow$	Light. $\uparrow$	Src.Int. $\uparrow$
Baseline	5.489 $\pm$ 1.48	0.811 $\pm$ 0.13	0.815 $\pm$ 0.05	0.835 $\pm$ 0.12	0.848 $\pm$ 0.11	0.754 $\pm$ 0.11	0.767 $\pm$ 0.10	0.902 $\pm$ 0.10
Direct	5.284 $\pm$ 1.54	0.867 $\pm$ 0.11	0.837 $\pm$ 0.04	0.870 $\pm$ 0.07	0.879 $\pm$ 0.07	0.782 $\pm$ 0.08	0.770 $\pm$ 0.09	0.906 $\pm$ 0.11
RaR (Gunjal et al., 2025)	5.533 $\pm$ 1.65	0.850 $\pm$ 0.14	0.829 $\pm$ 0.05	0.862 $\pm$ 0.08	0.863 $\pm$ 0.09	0.769 $\pm$ 0.09	0.770 $\pm$ 0.10	0.905 $\pm$ 0.13
IEC w/o GC	5.259 $\pm$ 1.51	0.890$\pm$0.11	0.852 $\pm$ 0.04	0.878 $\pm$ 0.07	0.887 $\pm$ 0.06	0.799 $\pm$ 0.07	0.799$\pm$0.08	0.918$\pm$0.10
IEC (Ours)	5.203$\pm$1.48	0.883$\pm$0.10	0.852$\pm$0.04	0.882$\pm$0.07	0.893$\pm$0.06	0.803$\pm$0.07	0.796 $\pm$ 0.08	0.910 $\pm$ 0.11

well with human quality judgments: GenEval (Ghosh et al., 2023) counts compositional errors in text-to-image task, and TIFA (Hu et al., 2023) counts question-answering failures for text-image alignment.

Cross-Family Judge Evaluation. Because both our training (GPT-5-mini) and evaluation (GPT-5) judges come from the same model family, we additionally evaluate, as an ablation, with a frozen open-source judge, Qwen3.6-35B-A3B (with temperature set to 0), using the identical CEC, RaR, and Direct pipelines, see §5.

4. Results

We evaluate on MDressBench (*flat* and *non-flat*). Prior to evaluation, we validate CEC as a selection metric: Table 3 shows CEC matches the human top choice twice as often as Direct or RaR (Top-1: 0.60 vs. 0.30), supporting its use as the primary evaluation signal. We organize our results around three questions: Does error counting outperform rubric-based and direct rewards? (§4.1) Why does explicit counting fail while implicit counting succeeds? (§5) What role does group calibration play? (§5)

4.1. Error Counting Outperforms Rubrics and Direct Scoring

Tables 1 and 2 present results on MDressBench. All RL methods improve over the SFT baseline. **IEC vs. RaR.** IEC outperforms RaR on all eight metrics under both flat and non-flat references. On flat references, IEC achieves CEC 5.31 vs. 5.60 for RaR. On the more challenging non-flat subset, IEC exceeds RaR by 5.96% (CEC), 3.88% (RaR score), 2.32% (garment transfer), 3.48% (attribute preservation), 4.33% (realism), 3.3% (lighting), and 0.6% (source

Table 3. Human alignment of automated selection rules ($n = 25 \times 8$). CEC achieves $2 \times$ higher top-1 accuracy than alternatives, validating it as a reliable selection metric.

Selector	Top-1 $\uparrow$	Top-3 $\uparrow$	Top-5 $\uparrow$
Direct	0.30	0.80	1.0
RaR (Gunjal et al., 2025)	0.30	0.90	1.0
CEC (Ours)	0.60	1.0	1.0

integrity). This gap is notable because RaR must generate rubrics *without* an ideal reference to anchor criteria; rubrics become generic and miss prompt-specific failure modes, producing misaligned reward signals. **IEC vs. Direct Scoring.** Direct scoring reduces CEC from 6.02 (SFT) to 5.37 on flat references, however, IEC still outperforms on all metrics. Direct scoring compresses quality into a single scalar, losing the fine-grained error sensitivity that IEC’s aspect decomposition provides. The gap widens on non-flat references, where transfer errors are more varied and harder to capture with a direct score. **RaR underperforms Direct.** Strikingly, RaR scores worse than Direct on CEC (5.60 vs. 5.37 on flat; 5.53 vs. 5.28 on non-flat). This confirms that in the reference-free setting, attempting to construct rubrics without ideal answers can be worse than not using rubrics at all. Binary pass/fail judgments across many criteria compound stochastic judge errors, and generic criteria miss the specific failure modes of each input.

4.2. Garment Category Analysis

Figures 6 and 1, and Tables 8 and 9 show that IEC improves performance *across all garment types* rather than shifting errors between categories. The largest gain is on lowers (CEC 4.66 vs. 5.65 for SFT), while tops are the only case where Direct is slightly better on CEC. IEC attains the

best score in 20 of 24 (metric, category) combinations on flat references, indicating gains are not driven by a single category. On non-flat references, IEC wins 14/24 combinations, with gains concentrating on dresses (all 8 metrics) and tops (5/8). Figure 7 in Appendix B.3 summarizes this per-category structure: IEC achieves the largest area across all categories, with the most pronounced gains on Dresses. See Appendix B.5 for a qualitative discussion of reward hacking modes across reward designs.

4.3. Generalization to Standard Benchmarks

To assess whether IEC’s gains transfer beyond MDress-Bench, we evaluate on the paired test splits of VITON-HD and DressCode using perceptual metrics. Table 6 compares our RL post-trained models against six external SFT baselines (Chong et al., 2024; 2025; Zeng et al., 2024; Choi et al., 2024; Xu et al., 2024; Yang et al., 2025). IEC achieves the best LPIPS (0.047) and SSIM (0.892) on VITON-HD, and the best FID (3.610) and ties for best SSIM (0.927) on DressCode. Across both benchmarks, IEC attains the top score on 6 of 8 metrics despite using only 60 RL steps with no additional paired data. The consistent gains over external SFT models, trained on substantially larger paired datasets, validate that error-counting rewards provide an efficient path to quality improvement.

5. Ablations

Why Not Explicit Error Counting? A natural question is whether exposing error lists (EEC) would provide even finer-grained signal than implicit scoring. Table 4 and Figure 2 answer this decisively: EEC *fails to improve* CEC on flat references (6.03 vs. 6.02 SFT) and *degrades* on non-flat references (5.72 vs. 5.49 SFT). IEC yields large improvements in both settings. This behavior is consistent with the reward dynamics in Fig. 2. EEC produces high-variance rewards and degrades early. Explicit label enumeration is brittle, as semantically identical errors can be counted under multiple surface forms, injecting noise into within-group comparisons and destabilizing rankings. IEC, by contrast, asks the judge to implicitly aggregate errors into a calibrated scalar, yielding more stable preference signals. More broadly, in reference-free settings, error-based rewards from multimodal VLM judges are inherently susceptible to variance from linguistic and perceptual ambiguity (see Appendix B.6). Implicit scoring is therefore not merely a convenience; it is a stability-critical design choice for optimization.

Group Calibration Table 2 includes IEC w/o GC (group calibration removed). Without calibration, IEC still outperforms all baselines, confirming that the core error-counting signal is strong. However, group calibration improves CEC (5.203 vs. 5.259 on non-flat) by reducing prompt-to-prompt




Figure 4. Effect of Best-of- N with 12 generations. Higher N typically reduces reward diversity and performance.

scale variability, which stabilizes GRPO’s advantage estimates. The effect is more pronounced on harder prompts where reward magnitudes vary most. Note that calibration is effective for IEC precisely because the within-group ranking is reliable; for EEC, where stochastic enumeration corrupts the ranking itself (Appendix B.6), calibration does not provide equivalent benefit.

Robustness to Underlying Judge Model Because our training and evaluation judges share a model family, we ablate judge identity by evaluating the non-flat subset of the benchmark with a different and open-weight judge, as described in §3. Table 5 shows the relative performance across methods and judges is preserved: IEC outperforms Baseline, Direct, and RaR under *both* judges on all three metrics, and the IEC>Direct>RaR ordering holds cross-family. This indicates IEC’s gains are not an artifact of self-preference within a single judge family. Full per-category results are in Appendix B.8.

Effect of Best-of- N Sampling Best-of- N selects the best $N/2$ and worst $N/2$ within a group to increase diversity. We vary $N \in \{8, 12\}$ with 12 generations fixed (Figure 4). Increasing N generally reduces aspect scores, consistent with a quality-diversity trade-off. CEC responds asymmetrically: it slightly worsens on non-flat references (5.20→5.27) but improves on flat (5.31→4.80). We use $N=8$ as it increases diversity within groups and is robust across both reference types.

Qualitative Analysis In Figure 14, RaR-trained models blur garment details, reflected in poor realism scores. IEC preserves fine-grained textures and patterns. Figure 15 shows IEC can occasionally fail on source integrity when source and reference backgrounds differ, accounting for the smaller 0.6% improvement. This small gain is structural rather than a signal failure, because source integrity has the smallest within-group reward variance of any axis at every training window ($\sigma_{\text{src}}=0.060-0.074$ vs. $\sigma_{\text{garm}}=0.102-0.170$), so GRPO’s group-normalized advantage on this axis is correspondingly weaker. The residual failures are predominantly pose/limb edits (38% of labels) rather than identity/face

Table 4. Explicit vs. implicit error counting. EEC fails to improve over baseline; IEC achieves consistent gains.

Method	Flat			Non-Flat		
	CEC↓	Rub.↑	Dir.↑	CEC↓	Rub.↑	Dir.↑
Baseline (SFT)	6.02	0.76	0.80	5.49	0.81	0.82
EEC	6.03	0.76	0.81	5.72	0.77	0.81
IEC (Ours)	5.31	0.88	0.85	5.20	0.88	0.85

Table 5. Judge-family ablation on MDressBench (non-flat). Evaluation using a different judge shows that IEC dominates Direct/RaR across both judges, while absolute rubric values shift under the stricter Qwen judge, the relative ordering holds.

Method	GPT-5			Qwen3.6-35B-A3B		
	CEC↓	Rub↑	Dir↑	CEC↓	Rub↑	Dir↑
Baseline	5.489	0.811	0.815	7.114	0.087	0.879
Direct	5.284	0.867	0.837	5.550	0.107	0.919
RaR (Gunjal et al., 2025)	5.533	0.850	0.829	5.905	0.106	0.907
IEC (Ours)	5.203	0.883	0.852	5.377	0.110	0.919

changes (30%). See Appendix B.5 for details.

Training Dynamics Figure 2 shows IEC yields smooth, monotonic reward improvement, while Direct and RaR are noisier and EEC regresses early. Figure 5 compares IEC and RaR across iterations (20, 40, 60): IEC consistently achieves lower CEC, reaching minimum of 4.80 at iteration 40 (20% below baseline) vs. 5.06 for RaR. Both methods regress slightly after iteration 40, consistent with reward overfitting observed in prior RL post-training work (Liu et al., 2025; Wang et al., 2025a). At iteration-40, with IEC a typical generation carries a tight error distribution (median 5, p90 6, max 10) with *no* zero-error candidates, whereas EEC reports zero errors on 24.5% of candidates while reaching 26 on others, direct evidence of the enumeration variance behind EEC’s instability. We report at iteration-60 as RaR starts to regress; over the 0-60 window IEC’s calibrated reward stays above its starting value ($\Delta_{60}=+0.018$), whereas RaR peaks around iteration-40 and is already declining by iteration-60, with its per-rubric-item fail rate climbing from 3% to 15% respectively. See Appendix B.7 for the full per-axis analysis and training dynamics beyond iteration 60.




Figure 5. Training dynamics. CEC ↓ for IEC and RaR across post-training iterations. IEC consistently achieves lower error.

6. Related Work

RL for Visual Generation. RLHF is standard for aligning language models (Ouyang et al., 2022); DPO (Rafailov et al., 2023) removes the explicit reward model, while GRPO (Shao et al., 2024) normalizes advantages within candidate groups. Several works adapt policy optimization to diffusion models (Black et al., 2023; Fan et al., 2024; Wallace et al., 2024). Flow-GRPO (Liu et al., 2025) introduces ODE-to-SDE conversion for flow-matching models; MixGRPO (Li et al., 2025a) uses mixed ODE-SDE sampling, SuperFlow (Chen et al., 2025) enables on-the-fly training, and TreeGRPO (Ding & Ye, 2025) constructs search trees for step-level credit assignment. DiffusionNFT (Zheng et al., 2025) proposes a negative-aware finetuning objective. GRPO-Guard (Wang et al., 2025a) addresses over-optimization via regulated clipping, directly relevant to the reward regression we observe (Fig. 5). Our work focuses on reward design rather than optimization mechanics.

Reward Design for Subjective Tasks. ImageReward and HPSv2 (Xu et al., 2023; Wu et al., 2023) train reward models on human preference datasets but are aesthetics-centric. Using VLMs as evaluators (Chen et al., 2024; Gong et al., 2025; Wu et al., 2024; Zhang et al., 2024; Wang et al., 2025b) has been applied to both evaluation and training-time rewards. RaR (Gunjal et al., 2025) generates task-specific rubrics, improving interpretability; Rubicon (Huang et al., 2025) extends this with defense rubrics, exploring explicit vs. implicit aggregation; and DR Tulu (Shao et al., 2026) evolves rubrics over training to counter generic, trivially-satisfied criteria, the same dilution we observe for reference-free rubrics. QA-LIGN (Jacob Dineen et al., 2025) decomposes a constitution into hierarchical principle-specific QA checks scored by an LLM judge and optimized with GRPO for safety alignment using a structured per-axis reward complementary to our error-counting method. UniWorld-V2/Edit-R1 (UniWorld Team, 2025) uses VLM implicit feedback for

Table 6. We evaluate SFT-only models trained on task-specific paired datasets vs. our RL post-training methods. Best in **bold**, second best underlined. IEC matches or surpasses all external baselines on 6/8 metrics despite using only 60 RL steps with no additional paired data.

Method	VITON-HD				DressCode			
	LPIPS↓	SSIM↑	FID↓	KID↓	LPIPS↓	SSIM↑	FID↓	KID↓
<i>SFT Methods</i>								
CatVTON (Chong et al., 2024)	0.057	0.870	5.430	0.410	0.046	0.892	3.990	0.820
CatVT2ON (Chong et al., 2024)	0.057	0.890	8.100	2.250	0.037	<u>0.922</u>	5.720	2.340
CAT-DM (Zeng et al., 2024)	0.080	<u>0.877</u>	5.600	0.830	—	—	—	—
IDM-VTON (Choi et al., 2024)	0.102	0.870	6.290	—	0.062	0.920	8.640	0.900
OOTDiffusion (Xu et al., 2024)	0.071	0.878	8.810	0.820	0.045	0.927	4.200	0.370
OmniVTON (Yang et al., 2025)	0.145	0.832	7.760	—	0.119	0.865	5.340	—
<i>Ours (SFT + RL Post-Training)</i>								
Baseline (SFT)	0.059	0.872	5.490	0.820	0.038	0.893	4.130	0.990
Direct	0.049	0.888	4.980	0.650	0.029	0.921	3.680	0.840
RaR (Gunjal et al., 2025)	<u>0.055</u>	0.878	5.210	<u>0.690</u>	0.031	0.915	<u>3.920</u>	0.940
IEC w/o GC	0.052	0.875	5.190	0.660	<u>0.034</u>	0.913	3.980	0.930
IEC (Ours)	0.047	0.892	<u>5.120</u>	0.660	0.035	0.927	3.610	<u>0.790</u>

image editing post-training via logit-based scoring, which is structurally parallel to our implicit approach, though applied to general editing. CoCoEdit (Wu et al., 2026) shows exclusive reliance on VLM rewards degrades content consistency, motivating region-based regularization—paralleling our observation that source integrity sees the smallest IEC improvement. GenEval (Ghosh et al., 2023) and TIFA (Hu et al., 2023) demonstrate that decomposing quality into countable criteria (compositional errors and QA failures, respectively) yields more stable signals than holistic scores. IEC extends this principle from evaluation to training rewards in a reference-free setting.

Virtual Try-On and Evaluation. VTO synthesizes images of a person wearing a target garment while preserving identity and background. Early methods used geometric warping (Han et al., 2018; Wang et al., 2018); VITON-HD (Choi et al., 2021) introduced high-resolution try-on, DressCode (Morelli et al., 2022) extended to multiple categories, LADI-VTON (Morelli et al., 2023) introduced latent diffusion with textual inversion, and TryOnDiffusion (Zhu et al., 2023) proposed parallel U-Nets. Recent methods explore DiT (Peebles & Xie, 2023) and flow-matching architectures, but rely on supervised fine-tuning on narrow datasets. VTON-VLLM (Wan et al., 2025) is the closest prior work: it aligns VTO via a VLLM fine-tuned on human-validated feedback, whereas we perform on-policy GRPO post-training using error-counting rewards without human preference labels or a trained reward model. Traditional VTO metrics (FID (Heusel et al., 2017), LPIPS (Zhang et al., 2018), SSIM (Wang et al., 2004)) correlate poorly with human perception of VTO-specific quality (Morelli et al., 2022) and cannot detect semantic errors like incorrect sleeve length. Our CEC metric addresses this via a shared error vocabulary for consistent, fine-grained measurement.

7. Limitations and Broader Considerations

Domain breadth. We validate IEC only on VTO; testing on other reference-free domains (image editing, robotic

manipulation, creative writing) would establish generality. **Missing baselines.** We do not compare against VTON-VLLM (Wan et al., 2025), which requires a fine-tuned VLLM reward model, whereas our method is training-free. **Metric validity.** CEC improves human alignment in our 25-sample study but remains an automated proxy; larger studies evaluating IEC outputs directly are future work. **Reproducibility.** IEC relies on non-deterministic LLMs as judges (GPT-5-mini/GPT-5); API changes could silently alter reward signals, motivating future work with open-source judges. **Fairness.** Our benchmark is curated from DressCode, which does not fully cover the diversity of body shapes, skin tones, and garment styles, and reward designs optimizing for judge-perceived quality may inherit judge biases. See Appendix B.4 for expanded discussion.

8. Conclusion

We identified the reference-free setting as a gap in RL post-training: domains where ideal answers are unavailable, multiple outputs are valid, and quality is defined by the absence of errors. We proposed Implicit Error Counting (IEC), showing that implicit score emission and group calibration are necessary for stable optimization. On VTO, IEC consistently outperforms rubric-based and direct-scoring rewards and matches or exceeds supervised baselines trained on far more paired data, at half the judge compute of RaR. When you cannot define what an ideal output looks like, define what a bad one looks like, and count.

References

- Black, K., Janner, M., Du, Y., Kostrikov, I., and Levine, S. Training diffusion models with reinforcement learning. *arXiv preprint arXiv:2305.13301*, 2023.
- Chen, D., Chen, R., Zhang, S., Wang, Y., Liu, Y., Zhou, H., Zhang, Q., Wan, Y., Zhou, P., and Sun, L. Mllm-as-a-judge: Assessing multimodal llm-as-a-judge with vision-language benchmark. In *Forty-first International*

- Conference on Machine Learning, 2024.
- Chen, K., Xu, Z., Shen, Y., Lin, Z., Yao, Y., and Huang, L. Superflow: Training flow matching models with rl on the fly. *arXiv preprint arXiv:2512.17951*, 2025.
- Choi, S., Park, S., Lee, M., and Choo, J. Viton-hd: High-resolution virtual try-on via misalignment-aware normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14131–14140, 2021.
- Choi, Y., Kwak, S., Lee, K., Choi, H., and Shin, J. Improving diffusion models for authentic virtual try-on in the wild. In *European Conference on Computer Vision*, pp. 206–235. Springer, 2024.
- Chong, Z., Dong, X., Li, H., Zhang, S., Zhang, W., Zhang, X., Zhao, H., Jiang, D., and Liang, X. Catvton: Concatenation is all you need for virtual try-on with diffusion models. *arXiv preprint arXiv:2407.15886*, 2024.
- Chong, Z., Zhang, W., Zhang, S., Zheng, J., Dong, X., Li, H., Wu, Y., Jiang, D., and Liang, X. Catv2ton: Taming diffusion transformers for vision-based virtual try-on with temporal concatenation. *arXiv preprint arXiv:2501.11325*, 2025.
- Ding, Z. and Ye, W. Treegrpo: Tree-advantage grpo for online rl post-training of diffusion models. *arXiv preprint arXiv:2512.08153*, 2025.
- Fan, Y., Watkins, O., Du, Y., Liu, H., Ryu, M., Boutilier, C., Abbeel, P., Ghavamzadeh, M., Lee, K., and Lee, K. Reinforcement learning for fine-tuning text-to-image diffusion models. In *Advances in Neural Information Processing Systems*, 2024.
- Ghosh, D., Hajishirzi, H., and Schmidt, L. Geneval: An object-focused framework for evaluating text-to-image alignment. *Advances in Neural Information Processing Systems*, 36:52132–52152, 2023.
- Gong, Y., Wang, X., Wu, J., Wang, S., Wang, Y., and Wu, X. Onereward: Unified mask-guided image generation via multi-task human preference learning. *arXiv preprint arXiv:2508.21066*, 2025.
- Gunjal, A., Wang, A., Lau, E., Nath, V., He, Y., Liu, B., and Hendryx, S. Rubrics as rewards: Reinforcement learning beyond verifiable domains. *arXiv preprint arXiv:2507.17746*, 2025.
- Han, X., Wu, Z., Wu, Z., Yu, R., and Davis, L. S. Viton: An image-based virtual try-on network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 7543–7552, 2018.
- Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., and Hochreiter, S. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, volume 30, 2017.
- Hu, Y., Liu, B., Kasai, J., Wang, Y., Ostendorf, M., Krishna, R., and Smith, N. A. Tifa: Accurate and interpretable text-to-image faithfulness evaluation with question answering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 20406–20417, 2023.
- Huang, Z., Zhuang, Y., Lu, G., Qin, Z., Xu, H., Zhao, T., Peng, R., Hu, J., Shen, Z., Hu, X., et al. Reinforcement learning with rubric anchors. *arXiv preprint arXiv:2508.12790*, 2025.
- Jacob Dineen, A. R., Liu, Q., Xu, Z., Ye, X., Shen, M., Li, Z., Lu, S., Baral, C., Chen, M., and Zhou, B. Qa-align: Aligning llms through constitutionally decomposed qa. *arXiv preprint arXiv:2506.08123*, 2025.
- Labs, B. F., Batifol, S., Blattmann, A., Boesel, F., Consul, S., Diagne, C., Dockhorn, T., English, J., English, Z., Esser, P., et al. Flux. 1 kontext: Flow matching for in-context image generation and editing in latent space. *arXiv preprint arXiv:2506.15742*, 2025.
- Li, J., Cui, Y., Huang, T., Ma, Y., Fan, C., Yang, M., and Zhong, Z. Mixgrpo: Unlocking flow-based grpo efficiency with mixed ode-sde. *arXiv preprint arXiv:2507.21802*, 2025a.
- Li, Q., Qiu, S., Han, J., Xu, X., Seyfioglu, M. S., Koo, K. K., and Bouyarmane, K. Dit-vton: Diffusion transformer framework for unified multi-category virtual try-on and virtual try-all with integrated image editing. *arXiv preprint arXiv:2510.04797*, 2025b.
- Lipman, Y., Chen, R. T., Ben-Hamu, H., Nickel, M., and Le, M. Flow matching for generative modeling. *arXiv preprint arXiv:2210.02747*, 2022.
- Liu, J., Liu, G., Liang, J., Li, Y., Liu, J., Wang, X., Wan, P., Zhang, D., and Ouyang, W. Flow-grpo: Training flow matching models via online rl. *arXiv preprint arXiv:2505.05470*, 2025.
- Liu, X., Gong, C., and Liu, Q. Flow straight and fast: Learning to generate and transfer data with rectified flow. *The International Conference on Learning Representations*, 2022.
- Morelli, D., Fincato, M., Cornia, M., Landi, F., Cesari, F., and Cucchiara, R. Dress code: High-resolution multi-category virtual try-on. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2231–2235, 2022.

- Morelli, D., Baldrati, A., Cartella, G., Cornia, M., Bertini, M., and Cucchiara, R. Ladi-vton: Latent diffusion textual-inversion enhanced virtual try-on. In *Proceedings of the ACM International Conference on Multimedia*, pp. 8580–8589, 2023.
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. Training language models to follow instructions with human feedback. *Advances in Neural Information Processing Systems*, 35:27730–27744, 2022.
- Peebles, W. and Xie, S. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4195–4205, 2023.
- Rafailov, R., Sharma, A., Mitchell, E., Manning, C. D., Ermon, S., and Finn, C. Direct preference optimization: Your language model is secretly a reward model. *Advances in Neural Information Processing Systems*, 36: 53728–53741, 2023.
- Seyfioglu, M. S., Bouyarmane, K., Kumar, S., Tavanaei, A., and Tutar, I. B. Diffuse to choose: Enriching image conditioned inpainting in latent diffusion models for virtual try-all. *arXiv preprint arXiv:2401.13795*, 2024.
- Shao, R., Asai, A., Shen, S. Z., Ivison, H., Kishore, V., Zhuo, J., Zhao, X., Park, M., Finlayson, S. G., Sontag, D., Murray, T., Min, S., Dasigi, P., Soldaini, L., Brahman, F., tau Yih, W., Wu, T., Zettlemoyer, L., Kim, Y., Hajishirzi, H., and Koh, P. W. Dr tulu: Reinforcement learning with evolving rubrics for deep research, 2026. URL <https://arxiv.org/abs/2511.19399>. Accessed: 2026-06-30.
- Shao, Z., Wang, P., Zhu, Q., Xu, R., Song, J., Bi, X., Zhang, H., Zhang, M., Li, Y., Wu, Y., et al. Deepseekmath: Pushing the limits of mathematical reasoning in open language models. *arXiv preprint arXiv:2402.03300*, 2024.
- UniWorld Team. UN I W O R L D-V2: Reinforce Image Editing with Diffusion Negative-Aware Finetuning and MLLM Implicit Feedback. Technical report, Peking University, Rabbitpre AI, 2025. Technical Report, arXiv:2510.16888.
- Wallace, B., Dang, M., Rafailov, R., Zhou, L., Lou, A., Purber, S., Ermon, S., Xiong, C., Joty, S., and Naik, N. Diffusion model alignment using direct preference optimization. *arXiv preprint arXiv:2311.12908*, 2024.
- Wan, S., Chen, J., Cai, Q., Pan, Y., Yao, T., and Mei, T. VTon-VLLM: Aligning virtual try-on models with human preferences. In *NeurIPS*, 2025.
- Wang, B., Zheng, H., Liang, X., Chen, Y., Lin, L., and Yang, M. Toward characteristic-preserving image-based virtual try-on network. In *Proceedings of the European Conference on Computer Vision*, pp. 589–604, 2018.
- Wang, J., Liang, J., Liu, J., Liu, H., Liu, G., Zheng, J., Pang, W., Ma, A., Xie, Z., Wang, X., Wang, M., Wan, P., and Liang, X. Grpo-guard: Mitigating implicit over-optimization in flow matching via regulated clipping. *arXiv preprint arXiv:2510.22319*, 2025a.
- Wang, Y., Zang, Y., Li, H., Jin, C., and Wang, J. Unified reward model for multimodal understanding and generation. *arXiv preprint arXiv:2503.05236*, 2025b.
- Wang, Z., Bovik, A. C., Sheikh, H. R., and Simoncelli, E. P. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- Wu, C., Li, J., Zhou, J., Lin, J., Gao, K., Yan, K., Yin, S.-m., Bai, S., Xu, X., Chen, Y., et al. Qwen-image technical report. *arXiv preprint arXiv:2508.02324*, 2025.
- Wu, H., Zhang, Z., Zhang, E., Chen, C., Liao, L., Wang, A., Li, C., Sun, W., Yan, Q., Zhai, G., et al. Q-bench: A benchmark for general-purpose foundation models on low-level vision. In *ICLR*, 2024.
- Wu, X., Hao, Y., Sun, K., Chen, Y., Zhu, F., Zhao, R., and Li, H. Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv preprint arXiv:2306.09341*, 2023.
- Wu, Y., Xie, C., Li, R., Chen, L., Yi, Q., and Zhang, L. Cooedit: Content-consistent image editing via region regularized reinforcement learning. *arXiv preprint arXiv:2602.14068*, 2026.
- Xu, J., Liu, X., Wu, Y., Wang, Y., Ye, W., Geng, S., Zhao, Y., Li, J., Li, C., Sun, H., et al. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36, 2023.
- Xu, Y., Gu, T., Chen, W., and Chen, C. Ootdiffusion: Outfitting fusion based latent diffusion for controllable virtual try-on, 2024. URL <https://arxiv.org/abs/2403.01779>. Accessed: 2026-06-30.
- Xue, Z., Wu, J., Gao, Y., Kong, F., Zhu, L., Chen, M., Liu, Z., Liu, W., Guo, Q., Huang, W., et al. Dancegrpo: Unleashing grpo on visual generation. *arXiv preprint arXiv:2505.07818*, 2025.
- Yang, Z., Li, Y., He, S., Li, X., Xu, Y., Dong, J., and Du, Y. Omnivton: Training-free universal virtual try-on. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 16702–16711, 2025.

- Zeng, J., Song, D., Nie, W., Tian, H., Wang, T., and Liu, A.-A. Cat-dm: Controllable accelerated virtual try-on with diffusion model. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 8372–8382, 2024.
- Zhang, R., Isola, P., Efros, A. A., Shechtman, E., and Wang, O. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 586–595, 2018.
- Zhang, Z., Wu, H., Zhang, E., Zhai, G., and Lin, W. Q-bench: A benchmark for multi-modal foundation models on low-level vision from single images to pairs. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024.
- Zheng, K., Chen, H., Ye, H., Wang, H., Zhang, Q., Jiang, K., Su, H., Ermon, S., Zhu, J., and Liu, M.-Y. Diffusion-nft: Online diffusion reinforcement with forward process. *arXiv preprint arXiv:2509.16117*, 2025.
- Zhu, L., Yang, D., Zhu, T., Reda, F., Chan, W., Saharia, C., Norouzi, M., and Kemelmacher-Shlizerman, I. Tryondiffusion: A tale of two unets. *arXiv preprint arXiv:2306.08276*, 2023.

A. Mismatch-Dress Benchmark

We construct the Mismatch-Dress Benchmark (MDressBench) through a VLM-based annotation pipeline designed to maximize attribute variation while enforcing strict visibility constraints.

Annotation Schema. We define a hierarchical garment taxonomy with three primary categories (upper, lower, dress), type-specific subcategories within each (18 upper, 13 lower, 13 dress types), and a structured attribute set characterizing each garment. Required attributes include sleeve length, fit, and garment length; optional attributes include neckline, pattern, rise, waistline, and color. Each attribute takes values from a discrete enumerated set (3–8 options), enabling deterministic annotation and programmatic filtering. We additionally annotate *garment presentation* $\in \{\text{masculine, feminine, unisex}\}$ independently of model gender—classification based on cut and silhouette rather than cultural convention—enabling analysis of try-on performance across gender-presentation boundaries.

Visibility and Extraction. A garment is included only if fully visible: critical boundaries (hems, cuffs, necklines, waistbands) must be unoccluded and within frame, with $\geq 90\%$ visibility required for layered garments. Accessories are explicitly excluded. Our VLM annotation uses two-stage prompting: a system prompt defining the ontology, visibility rules, and output schema, paired with a user prompt injecting the known primary category and maximum item count. Category-specific extraction rules prevent ambiguity: upper/lower images ($n_{\max} = 2$) include the ground truth category and may include the complementary category if fully visible, with coordinated outfits (e.g., shirt + skirt) split into components rather than mislabeled as dresses; dress images ($n_{\max} = 1$) enforce single-piece annotation.

Attribute Diversity. Let $\mathcal{S}$ denote the space of attribute combinations and $\hat{P}(s)$ the empirical frequency of combination s after initial annotation. We identify underrepresented strata as:

$$\mathcal{S}_{\text{rare}} = \left\{ s \in \mathcal{S} : \hat{P}(s) < \frac{0.5}{|\mathcal{S}|} \right\} \quad (7)$$

and actively source additional images to fill these gaps. This stratified resampling yields a balanced distribution satisfying $\max_s P(s) / \min_s P(s) \leq 5$. Allowing complementary garments in upper/lower images further increases attribute diversity while maintaining ground truth integrity. The resulting benchmark contains 700 challenging pairs spanning 44 garment types, achieving $\geq 60\%$ coverage of theoretically possible attribute combinations.

B. Training and Results

B.1. Training Hyperparameters

Table 7 summarizes the hyperparameters used for supervised fine-tuning (SFT) and reinforcement learning (RL) stages.

B.2. Results across MDressBench Garment Categories

Tables 8, 9 and Figure 6 provide complete per-category breakdowns for all metrics on flat and non-flat reference images respectively.

B.3. Garment-Category Radar Plot

Figure 7 shows per-category quality metrics on non-flat references. IEC consistently achieves the largest area across all categories, with the most pronounced gains on Dresses.

B.4. Expanded Limitations

Evaluation axes. We use five manually defined axes for VTO. Automating axis discovery for new domains—perhaps by having the judge itself propose relevant error categories—is an open direction.

Dynamic error weighting. We use equal axis weights throughout training. A curriculum that emphasizes structural correctness early and shifts toward fine-grained quality later may improve sample efficiency, analogous to curriculum-based rubric weighting suggested by (Gunjal et al., 2025).

Table 7. Training hyperparameters for SFT and RL stages.

Hyperparameter	SFT	RL
<i>Optimization</i>		
Batch size	4	3
Gradient accumulation steps	12	4
Learning rate	5e-5	1e-5
Weight decay	1e-4	1e-4
Max gradient norm	1.0	–
LR warmup steps	0	–
Mixed precision	bf16	bf16
<i>RL-specific</i>		
Num generations (K)	–	12
Best-of- n down-selection	–	8
Sampling steps	–	30
Resolution ($H \times W$)	512×512	512×512
CFG / guidance scale	–	2.5
Shift	–	3
Init same noise	–	✓
Timestep fraction	–	0.3
η (EMA)	–	0.5
Clip range	–	0.2
Advantage clip max	–	5.0
KL weight	–	0.1
Judge model	–	GPT-5-mini




Figure 6. Garment-category breakdown on MDressBench for flat (top) and non-flat (bottom) references. IEC (green) consistently outperforms baselines across all categories.

Benchmark coverage. MDressBench intentionally concentrates on hard, attribute-mismatched pairs and is curated from DressCode only, which does not cover all body shapes and garment styles. We focus on garments and ignore accessories like bags.

Table 8. Full per-category results for flat reference images. Best values per (metric, category) in **bold**.

Model	Cat	CEC↓	RaR↑	Direct↑	Garm.↑	Attr.↑	Real.↑	Light.↑	Src.Int.↑
Baseline	all	6.018±1.730	0.760±0.140	0.796±0.060	0.793±0.160	0.815±0.130	0.735±0.120	0.760±0.100	0.914±0.08
	top	5.220±1.150	0.753±0.140	0.801±0.030	0.812±0.100	0.818±0.130	0.726±0.100	0.737±0.090	0.936±0.03
	lower	5.652±1.250	0.777±0.100	0.776±0.030	0.795±0.120	0.768±0.140	0.707±0.120	0.734±0.110	0.900±0.10
	dress	6.072±1.760	0.760±0.140	0.796±0.060	0.792±0.170	0.816±0.130	0.736±0.120	0.762±0.100	0.913±0.08
Direct	all	5.373±1.570	0.854±0.110	0.840±0.040	0.864±0.090	0.873±0.080	0.785±0.080	0.771±0.090	0.925±0.07
	top	4.701 ±1.050	0.833 ±0.090	0.828±0.030	0.856±0.070	0.865±0.070	0.759±0.080	0.738±0.090	0.941±0.03
	lower	5.250±1.800	0.856±0.060	0.809±0.030	0.847 ±0.060	0.788±0.150	0.759 ±0.060	0.755±0.080	0.918±0.07
	dress	5.415±1.590	0.855±0.110	0.842±0.050	0.865±0.100	0.876±0.080	0.787±0.080	0.773±0.090	0.924±0.07
RaR	all	5.596±1.690	0.826±0.140	0.833±0.050	0.858±0.090	0.858±0.090	0.771±0.090	0.764±0.100	0.929±0.06
	top	5.066±1.070	0.805±0.110	0.829±0.030	0.856±0.070	0.865±0.070	0.761±0.070	0.740±0.090	0.943±0.03
	lower	5.286±1.220	0.824±0.070	0.799±0.040	0.833±0.080	0.772±0.140	0.747±0.080	0.749±0.100	0.921±0.09
	dress	5.634±1.720	0.827±0.140	0.834±0.050	0.859±0.090	0.859±0.090	0.772±0.090	0.765±0.100	0.929±0.06
EEC	all	6.032±1.860	0.761±0.190	0.807±0.060	0.846±0.100	0.849±0.090	0.735±0.100	0.709±0.110	0.915±0.06
	top	5.355±1.010	0.714±0.150	0.797±0.030	0.834±0.070	0.848±0.070	0.711±0.080	0.677±0.100	0.935±0.03
	lower	5.830±1.580	0.762±0.110	0.776±0.030	0.821±0.080	0.775±0.140	0.710±0.080	0.696±0.090	0.903±0.07
	dress	6.076±1.900	0.764±0.190	0.808±0.060	0.847±0.100	0.851±0.090	0.737±0.100	0.711±0.110	0.914±0.06
IEC	all	5.312 ±1.720	0.875 ±0.090	0.854 ±0.040	0.873 ±0.080	0.886 ±0.070	0.803 ±0.070	0.794 ±0.070	0.933 ±0.06
	top	4.951±1.250	0.819±0.120	0.838 ±0.030	0.861 ±0.060	0.867 ±0.080	0.774 ±0.070	0.766 ±0.080	0.945 ±0.03
	lower	4.661 ±1.790	0.872 ±0.040	0.813 ±0.030	0.839±0.070	0.802 ±0.150	0.759 ±0.070	0.761 ±0.080	0.929 ±0.05
	dress	5.347 ±1.740	0.878 ±0.090	0.856 ±0.040	0.875 ±0.090	0.889 ±0.070	0.806 ±0.070	0.796 ±0.070	0.932 ±0.06

 Table 9. Full per-category results for non-flat reference images. Best values per (metric, category) in **bold**.

Model	Cat	CEC↓	RaR↑	Direct↑	Garm.↑	Attr.↑	Real.↑	Light.↑	Src.Int.↑
Baseline	all	5.489±1.480	0.811±0.130	0.815±0.050	0.835±0.120	0.848±0.110	0.754±0.110	0.767±0.100	0.902±0.10
	top	5.049±1.030	0.837±0.070	0.808±0.030	0.828±0.090	0.845±0.100	0.741±0.090	0.742±0.100	0.922±0.08
	lower	6.089±1.980	0.759±0.180	0.743±0.080	0.758±0.130	0.818±0.110	0.655±0.140	0.680±0.150	0.846±0.15
	dress	5.502±1.480	0.811±0.130	0.817±0.050	0.837±0.120	0.848±0.110	0.757±0.110	0.770±0.100	0.902±0.10
Direct	all	5.284±1.540	0.867±0.110	0.837±0.040	0.870±0.070	0.879±0.070	0.782±0.080	0.770±0.090	0.906±0.11
	top	4.559 ±0.990	0.846 ±0.060	0.827±0.030	0.851±0.090	0.873±0.070	0.756±0.090	0.754±0.080	0.929±0.07
	lower	5.804±1.140	0.848±0.130	0.761±0.080	0.787±0.100	0.816±0.120	0.680±0.110	0.686±0.110	0.873±0.14
	dress	5.316±1.560	0.868±0.110	0.840±0.040	0.873±0.070	0.881±0.070	0.786±0.080	0.773±0.090	0.905±0.11
RaR	all	5.533±1.650	0.850±0.140	0.829±0.050	0.862±0.080	0.863±0.090	0.769±0.090	0.770±0.100	0.905±0.13
	top	4.704±1.270	0.838±0.120	0.827±0.040	0.845±0.090	0.868±0.070	0.756±0.090	0.753±0.090	0.940 ±0.04
	lower	6.286±1.560	0.852 ±0.130	0.783 ±0.080	0.798 ±0.110	0.838 ±0.090	0.694 ±0.120	0.724 ±0.100	0.898 ±0.11
	dress	5.565±1.650	0.851±0.140	0.830±0.050	0.864±0.080	0.863±0.090	0.772±0.090	0.772±0.100	0.903±0.13
EEC	all	5.721±1.700	0.769±0.180	0.806±0.050	0.856±0.080	0.855±0.080	0.739±0.090	0.712±0.110	0.893±0.12
	top	4.980±1.320	0.763±0.190	0.802±0.040	0.835±0.080	0.852±0.070	0.725±0.090	0.697±0.100	0.926±0.06
	lower	5.411 ±1.300	0.748±0.110	0.755±0.050	0.787±0.100	0.801±0.110	0.667±0.100	0.661±0.110	0.885±0.10
	dress	5.771±1.720	0.770±0.180	0.808±0.050	0.859±0.070	0.857±0.080	0.741±0.090	0.714±0.110	0.891±0.12
IEC	all	5.203 ±1.480	0.883 ±0.100	0.852 ±0.040	0.882 ±0.070	0.893 ±0.060	0.803 ±0.070	0.796 ±0.080	0.910 ±0.11
	top	5.079±1.250	0.826±0.080	0.837 ±0.030	0.854 ±0.090	0.878 ±0.080	0.772 ±0.070	0.772 ±0.080	0.938±0.05
	lower	6.071±1.570	0.841±0.150	0.773±0.080	0.798 ±0.100	0.827±0.100	0.691±0.110	0.710±0.110	0.869±0.15
	dress	5.191 ±1.480	0.888 ±0.090	0.855 ±0.040	0.885 ±0.070	0.895 ±0.060	0.807 ±0.070	0.800 ±0.080	0.909 ±0.11

B.5. Reward Hacking Modes

Reward hacking manifests differently across reward designs. Based on qualitative inspection of training outputs:

RaR. The most visible hacking mode is *texture smoothing*: RaR-trained models tend to blur garment details (Figure 14), likely because the binary pass/fail rubric criteria are easier to satisfy with smooth outputs that avoid fine-grained pattern errors. This manifests as washed-out textures and reduced pattern fidelity.

Direct. Direct scoring tends to produce outputs with increased color saturation and contrast, visually striking images that score well on a holistic quality scale but may not faithfully reproduce the reference garment’s actual appearance. Figure 8




Figure 7. **Garment-category quality metrics** on **non-flat** references. IEC (green) consistently achieves the largest area across all categories, with the most gains on Dresses.

shows qualitative comparisons on flat references of this issue.

IEC. IEC’s per-axis decomposition appears to reduce but not eliminate hacking. The primary failure mode is background alteration (Figure 15): when source and reference backgrounds differ substantially, the model sometimes modifies the source background to better match the reference, reducing the source integrity score. This is reflected in the smallest improvement margin on the source integrity axis (0.6% over RaR).

EEC. EEC exhibits the most severe instability. The explicit error lists fluctuate dramatically between near-identical outputs, causing the model to chase inconsistent reward signals rather than systematically improving quality. This manifests as early training regression (Figure 2).

B.6. Case Study: Within-Group Reward Comparison

Figure 9 shows the same source–reference pair (a brown sheer tulle dress with long sleeves over a white underdress) evaluated by IEC, EEC, and RaR during training. For each method, 8 candidate generations are ranked by reward from highest (A) to lowest (H). The reference dress has distinctive features: long sheer sleeves, a high neckline with button/collar detailing, and a floor-length layered skirt. This case study exposes concrete failure modes in EEC and RaR reward signals.

EEC produces rank inversions from stochastic enumeration. EEC assigns generation A a perfect reward of 1.000 by returning *zero errors* despite A visibly omitting the long sleeves that are a defining feature of the reference garment. Generation B, which correctly transfers the sleeves, receives 0.874 (2 errors: “flattened tulle skirt” and “visible seam artifact at waist”). The judge simply failed to enumerate the sleeve omission for A but caught texture errors in B, producing a rank inversion where a worse generation is rewarded higher. The pattern extends further: F appears among the best generations in the group but is relegated to 6th place (0.308, 11 errors). At the extremes, G and H receive 0.057 and 0.000 (15 and 17 errors respectively). The error count range of 0–17 across this single group illustrates the scale of stochastic variance.

Group calibration (Eqs. 3–4 in the main paper) adjusts reward *magnitudes* across prompts via a monotone transform but preserves within-group ordering. Since EEC’s ranking is itself corrupted ($A > B$ despite B being the better generation), calibration propagates the rank inversion unchanged into GRPO advantage estimates: A receives positive advantage and F



Figure 8. **Qualitative comparison of reward designs on flat references:** Given the pose-stitched masked source and the flat reference, IEC avoids *attribute preservation* errors and *garment transfer* failures over all other reward designs.

receives a negative advantage, directly reinforcing the wrong behavior. Similarly, CEC’s pooling phase was designed to address exactly this class of failure at evaluation time: by building a shared error vocabulary and re-evaluating candidates against it, errors missed for one candidate (such as “missing long sleeves”) are caught on re-examination because the pool contains sleeve-related errors from other candidates.

RaR misranks due to generic rubric criteria. RaR ranks A (0.855) and B (0.842) highest, yet both exhibit neckline misalignment and chest-area attribute preservation failures visible on inspection. Meanwhile, generations that better preserve the reference garment’s structure (e.g., C at 0.724 or G at 0.329) are ranked lower. The rubric contains 37 items, of which 19 are generic artifact checks (R19–R29: background corruption, ghosting, tiling, warping, halos, limb duplication, blur, banding, brush strokes, occlusion order, shadow mismatch) that all generations pass trivially. These items dilute the reward signal: passing 19 “free” items inflates scores for generations with garment-specific failures, while the few rubric items that do capture transfer fidelity (R6: hem style, R8: texture, R9: hardware) are weighted equally with the generic checks. Binary pass/fail scoring compounds the problem—a generation with a subtle neckline shift and one with a catastrophic sleeve omission both simply “fail” the relevant item with the same penalty.

IEC produces stable, interpretable rankings. IEC cleanly separates the group into generations that correctly transfer the dress (A–D, rewards 0.679–0.852) and those that omit the sleeves (E–H, rewards 0.199–0.371). The aspect decomposition is interpretable: A–D receive garment_transfer scores of 0.90 with only minor errors noted (“minor misalignment of small bodice details,” “slight ghosting through sheer areas”), while E–H drop to garment_transfer 0.60–0.65 with the judge consistently identifying “Long sleeves from reference are missing.” The reward gradient within each tier is smooth, providing GRPO with a reliable signal for advantage computation. Critically, IEC avoids the rank inversions seen in EEC because the judge internally aggregates error severity into calibrated scores—even if the textual error summaries vary in phrasing, the resulting scalar rewards remain consistent.

B.7. Training-time Failure Analysis

We generalize that case study quantitatively from the full IEC training logs, and place IEC’s behavior alongside RaR, EEC, and Direct. We use two analysis windows: the *iter-40 peak* ($n=216$ candidates over steps 36–44) and *iter-60* ($n=120$ over steps 56–60). The full 0–60 trajectory is shown for evolution analyses. A consistency check: per-axis training-judge means at iter-40 (garment 0.881, attribute 0.896, realism 0.806, lighting 0.794, source 0.940) match the evaluation numbers in Tab. 2 within ± 0.03 on every axis

Per-image error count at the reported checkpoint. Figure 10 plots the per-image error count distribution for IEC and EEC at the iter-40 peak. IEC’s distribution is concentrated (median 5, mean 5.02, p90 6, max 10), and *no candidate has zero errors*. EEC’s distribution is bimodal: 24.5% of candidates are reported with zero errors, while the remaining tail reaches p90=15 and max=26, the variance-of-enumeration behavior of EEC discussed in § 5 and Supp. B.6. Notably the EEC zero-error fraction is *highest at the iter-40 peak* (24.5%, vs 12%–16% in early training and 14.2% at iter-60).




Figure 9. **Within-group reward comparison across three reward designs** on the same source-reference pair (brown sheer tulle dress with long sleeves). Each row shows 8 generations ranked by reward, split into top-4 and bottom-4. **EEC** assigns generation A a perfect 1.000 (zero errors detected) despite A omitting the reference sleeves; B correctly transfers sleeves but receives 0.874. Error counts range from 0 to 17 across the group. **RaR** ranks A (0.855) and B (0.842) highest despite neckline/chest attribute failures, because 19 of 37 rubric items are generic artifact checks that all generations pass trivially, diluting garment-specific signal. **IEC** cleanly separates generations with correct sleeve transfer (A–D, 0.679–0.852) from those missing sleeves (E–H, 0.199–0.371), with interpretable aspect scores (garment.transfer drops from 0.90 to 0.60–0.65 when sleeves are absent). Group calibration can correct IEC’s prompt-to-prompt scale drift but cannot fix EEC’s rank inversions, as calibration preserves within-group ordering.

How the failure distribution evolves through training. Figure 11 shows the per-axis IEC training reward across iterations 1–60 with EMA smoothing ($\alpha=0.99$). All five axes peak together near iter 40; *garment transfer* and *attribute preservation* carry the largest improvement to peak (+0.110 and +0.090 over baseline). *Realism* rises to +0.087 at peak then *regresses to baseline* by iter-60. *Lighting* peaks at +0.036 and is already −0.015 below baseline at iter-60. *Source integrity* barely moves (+0.022 at peak), partly because its baseline (0.918) already sits near the per-axis ceiling. The *kind* of error the judge enumerates also evolves. Clustering the IEC error labels into 11 keyword themes, the dominant theme shifts from *fit/silhouette* (24% at steps 1-5) to *boundary/seam* (24% at iter-40 peak; *drape/geometry* also rises). Gross silhouette / fit mistakes are resolved first; finer boundary, drape, and blending mistakes persist later. The per-axis error-theme map at iter-40 is highly interpretable and aligns with each axis’s intended definition: *garment transfer* errors are dominated by *fit/silhouette* (48%); *attribute preservation* by *texture/material* (46%); *realism* by *boundary/seam* (53%); *lighting* by *lighting/shadow* (94%); *source integrity* by *pose/limb* (38%).

Source-integrity bottleneck. The smallest improvement margin in main Tab. 2 is on Source Integrity (+0.6%). Three lines of evidence in the training logs show this is mechanism-driven, not signal-failure:

(a) *Within-group reward variance.* Figure 12 shows the mean within-group reward standard deviation per axis at three training windows over the 0–60 horizon. Source integrity is the smallest at every window ($\sigma=0.060\text{--}0.074$), $1.4\text{--}2.5\times$ smaller than garment transfer ($\sigma=0.102\text{--}0.170$). Because GRPO’s normalized advantage scales as $\hat{A}=\frac{R-\bar{R}}{\sigma_R}$, a structurally smaller within-group σ produces a structurally smaller advantage signal on this axis, and thus a smaller per-step learning.

(b) *Error-label vocabulary.* The dominant theme in source integrity error labels is *pose/limb* (38%; arm, leg, hand, occlusion, clutch, bag), *not identity/face* (30%) or *background/scene* (25%). Failures look like “arm pose restructured to fit reference long sleeve” or “bare legs replaced with pants” rather than “face altered”.

Lighting and source-integrity are mechanistically coupled. We observe that lighting errors mention *background/scene* keywords in 63% of labels (alongside the design-intended *lighting/shadow* at 94%). Combined with source integrity’s dominant theme of *pose/limb* and its 25% secondary theme of *background/scene*, this suggests that a background-alteration hacking mode does not just hurt source integrity, but also hurts lighting, because changing the scene context changes what counts as “lighting consistent with the source”.

Reward-design temporal stability. Figure 13 shows the calibrated reward trajectories of IEC, RaR, and Direct over their full logged range. Using Δ_{60} (calibrated reward at iter-60 minus the first-5-step baseline) and Δ_{end} (last logged step minus baseline), IEC improves smoothly and is the only design still above baseline at the end of the run ($\Delta_{60}=+0.018$,




Figure 10. Per-image error-count distribution at the iter-40. Panel A overlays IEC (concentrated, 0% zero) and EEC (bimodal, 24.5% zero plus a long tail). Panels B-F are per-axis IEC error-count histograms; vertical lines mark mean (dashed) and median (dotted).

$\Delta_{\text{end}}=+0.018$). RaR peaks near iter-40 ($\Delta_{60}=+0.044$) and then collapses to $\Delta_{\text{end}}=-0.160$, and Direct peaks early and drifts back to $\Delta_{\text{end}}=-0.014$. We report the iter-60 checkpoint precisely because it precedes this late regression. Reward magnitudes differ across designs because the reward functions differ; the within-design shape, monotone for IEC versus peak-then-collapse for RaR, is the comparable quantity.

The RaR collapse is mirrored at the *judge* level: the per-rubric-item fail rate (fraction of rubric items the judge marks failed) is 11% at steps 1–5 and 14% at 6–19, drops to 3% at the iter-40 reward peak (the judge passes nearly everything once the model has improved), then rebounds to 15% at iter-60 and 19% by end of run. The same judge thus gives different verdicts on similar-quality outputs at different points in training, evidence that RaR’s reward is non-stationary rather than merely noisy.

B.8. Cross-family Judge: Full per-category results

To address the judge family-correlation concern, we re-evaluated using the open-source judge Qwen3.6-35B-A3B with the identical CEC, RaR, and Direct pipelines, see Table 10, on the non-flat MDressBench pairs.

Three patterns are robust across the two judge families: (i) **all four RL post-trained models beat the SFT baseline by 1.5–2.0 CEC** (Qwen: 5.09–5.91 vs 7.11), (ii) **IEC dominates Direct and RaR on every aggregate metric under both judges** (Qwen: IEC CEC 5.377 vs Direct 5.550, RaR 5.905; Rub IEC 0.110 vs Direct 0.107, RaR 0.106; Dir IEC 0.919 vs Direct 0.919, RaR 0.906), and (iii) **the IEC>Direct>RaR ordering is preserved cross-family** on all eight metrics including per-axis.

C. Group Relative Policy Optimization for Flow Models

C.1. Preliminaries

Flow Matching. Given a data sample $x_0 \sim X_0$ from the true distribution with condition c , and Gaussian noise $x_1 \sim \mathcal{N}(0, I)$, Rectified Flow (Liu et al., 2022; Lipman et al., 2022) defines interpolated samples as:

$$x_t = (1 - t)x_0 + tx_1, \quad t \in [0, 1]. \quad (8)$$

Under this convention, $t = 0$ corresponds to data and $t = 1$ to noise; generation proceeds by solving the ODE backward from $t = 1$ to $t = 0$. A neural network $v_\theta(x_t, t, c)$ is trained to approximate the velocity field $v = x_1 - x_0$ by minimizing:

$$\mathcal{L}_{FM}(\theta) = \mathbb{E}_{t, x_0, x_1} [\|v - v_\theta(x_t, t, c)\|_2^2]. \quad (9)$$




Figure 11. Per-axis IEC training reward trajectory. Faded: raw mean across the 24 candidates per step. Solid: EMA $\alpha=0.99$. Dashed verticals: paper iterations 20, 40, 60. Bottom-right: aggregate over all five axes.

Table 10. Full per-category results under the open-source judge on MDressBench **non-flat** references. IEC remains best-or-near-best across metrics

Model	Cat	CEC↓	Rub↑	Dir↑	Garm.↑	Attr.↑	Real.↑	Light.↑	Src.Int.↑
Baseline	all	7.114	0.087	0.879	0.878	0.895	0.868	0.885	0.918
	top	7.444	0.211	0.868	0.868	0.897	0.860	0.871	0.931
	lower	6.714	0.045	0.829	0.837	0.865	0.803	0.834	0.875
	dress	7.103	0.081	0.880	0.879	0.895	0.870	0.887	0.919
Direct	all	5.550	0.107	0.919	0.919	0.927	0.913	0.920	0.940
	top	6.773	0.201	0.882	0.883	0.909	0.869	0.883	0.932
	lower	6.080	0.042	0.846	0.854	0.868	0.832	0.846	0.891
	dress	5.467	0.103	0.923	0.923	0.929	0.918	0.924	0.941
RaR	all	5.905	0.106	0.906	0.907	0.917	0.899	0.909	0.932
	top	6.444	0.191	0.881	0.884	0.914	0.866	0.882	0.939
	lower	4.714	0.036	0.892	0.899	0.917	0.874	0.889	0.941
	dress	5.899	0.102	0.908	0.908	0.917	0.902	0.911	0.931
IEC	all	5.377	0.110	0.919	0.919	0.927	0.913	0.921	0.941
	top	6.618	0.168	0.896	0.896	0.916	0.887	0.898	0.951
	lower	5.625	0.027	0.868	0.876	0.891	0.851	0.863	0.911
	dress	5.299	0.108	0.922	0.922	0.929	0.916	0.923	0.941

C.2. GRPO for Flow Models

For each condition c , the flow model samples K images $\{x_0^i\}_{i=1}^K$ via reverse-time trajectories $\{(x_T^i, x_{T-1}^i, \dots, x_0^i)\}_{i=1}^K$, where T is the number of discretization steps. GRPO optimizes:

$$\mathcal{J}_{\text{GRPO}}(\theta) = \mathbb{E}_{c, \{x^i\}_{i=1}^K \sim \pi_{\theta_{\text{old}}}(\cdot|c)} \left[\frac{1}{K} \sum_{i=1}^K \frac{1}{T} \sum_{t=1}^T L_t^i \right], \quad (10)$$

where the per-step loss is:

$$L_t^i = \min \left(r_t^i \hat{A}^i, \text{clip}(r_t^i, 1-\epsilon, 1+\epsilon) \hat{A}^i \right) - \beta D_{KL}(\pi_{\theta} \parallel \pi_{\text{ref}}), \quad (11)$$

with probability ratio:

$$r_t^i = \frac{p_{\theta}(x_{t-1}^i | x_t^i, c)}{p_{\theta_{\text{old}}}(x_{t-1}^i | x_t^i, c)}. \quad (12)$$




Figure 12. Mean within-group reward σ per axis at three training windows ($\leq \text{iter-60}$). Source integrity is the smallest at every window. Smaller $\sigma \Rightarrow$ smaller GRPO advantage on that axis.




Figure 13. Cross-design calibrated reward (`reward_mean` per step). Left three panels: per-design raw and EMA ($\alpha=0.99$). Right: EMA of all three on common axes. Dashed verticals mark iterations 20/40/60. IEC keeps improving while RaR and Direct peak early and decline.

The advantage $\hat{A}^i$ is computed from the final generated image and applied uniformly across all timesteps:

$$\hat{A}^i = \frac{R(x_0^i, c) - \text{mean}_j (R(x_0^j, c))}{\text{std}_j (R(x_0^j, c))}. \quad (13)$$

C.3. ODE to SDE Conversion

To enable stochastic sampling required for the probability ratio, we convert the deterministic ODE $dx_t = v_\theta dt$ into an equivalent SDE that preserves marginal distributions (Liu et al., 2025):

$$x_{t+\Delta t} = x_t + v_\theta(x_t, t)\Delta t + \frac{\sigma_t^2}{2t} (x_t + (1-t)v_\theta(x_t, t)) \Delta t + \sigma_t \sqrt{\Delta t} \epsilon, \quad (14)$$

where $\epsilon \sim \mathcal{N}(0, I)$ and the noise schedule is $\sigma_t = a\sqrt{t/(1-t)}$ with $a = 1$ in all experiments.

The transition distribution is Gaussian: $p_\theta(x_{t-1}|x_t, c) = \mathcal{N}(\mu_\theta(x_t, t, c), \sigma_t^2 I)$, where μ_θ is the deterministic component (first two terms above). This yields a closed-form KL divergence:

$$D_{KL}(\pi_\theta \parallel \pi_{\text{ref}}) = \frac{\|\mu_\theta - \mu_{\text{ref}}\|^2}{2\sigma_t^2 \Delta t} = \frac{\Delta t}{2} \left(\frac{\sigma_t(1-t)}{2t} + \frac{1}{\sigma_t} \right)^2 \|v_\theta - v_{\text{ref}}\|^2. \quad (15)$$




Figure 14. **Qualitative analysis:** Post-training with RaR tends to blur out garment details as shown in these examples. The garment patterns in both dresses are maintained by IEC but visually faded with RaR.




Figure 15. **Qualitative analysis:** IEC tends to alter the background area when the source background is sanitized, and the reference background is not, which is reflected in the smallest, 0.6%, improvement over RaR on Source Integrity.

Remark. In our experiments, we set $\beta = 0$ (no KL regularization), as we found the clipped objective sufficient to prevent policy divergence over the short 60-step RL training horizon.

D. Training-Free Reward Design

D.1. Direct Reward

Our simplest reward baseline is a direct LLM-as-judge Likert score, analogous to the DIRECT-LIKERT baseline in RaR (Gunjajal et al., 2025). The judge is prompted to output a scalar

$$R_i^{\text{direct}} \in [0, 1] \quad (16)$$

describing overall try-on quality, together with per-aspect scores along the five axes above. In practice we use only the top-level reward field for GRPO training.

The system prompt decomposes evaluation with scoring bands: $[0.8, 1.0]$ for strong results with minor flaws, $[0.5, 0.8)$ for mixed quality, $[0.2, 0.5)$ for poor results with serious errors, and $[0.0, 0.2)$ for failed try-ons. The final reward may be biased downward if any aspect is catastrophically poor.

This reward is inexpensive (single pass per candidate) but aggregates all error modes into one unstructured scalar, which can make it harder to understand failure patterns and may lead to inconsistent scoring across diverse editing scenarios. See prompt in Figure 16.

D.2. Rubrics as a Reward

To obtain more structured supervision, we adapt Rubrics as Rewards (Gunjal et al., 2025) to virtual try-on. For each $(x_{\text{src}}, x_{\text{ref}})$ pair we first synthesize a prompt-specific rubric with 18–30 criteria:

$$\mathcal{R} = \{(\text{id}_j, \text{title}_j, \text{desc}_j, \text{cat}_j, w_j)\}_{j=1}^J, \quad (17)$$

where $\text{cat}_j \in \{\text{Essential}, \text{Important}, \text{Optional}, \text{Pitfall}\}$ denotes the criterion category and w_j is a signed weight (positive for the first three categories, negative for pitfalls).

Rubric Generation. Following the desiderata of (Gunjal et al., 2025), effective rubrics should be: (i) grounded in domain expertise, (ii) comprehensive in coverage, (iii) weighted by importance, and (iv) self-contained for independent evaluation. Our generation prompt instructs the VLM to produce items covering transfer fidelity, attribute preservation, realism, lighting, and source integrity. We mandate a “Global Artifact & Corruption” section with N pitfall items (background corruption, ghosting, tiling, warping, halos, duplicated limbs, blur, banding, occlusion errors, environment staining) phrased positively as “Avoids X...” with negative weights. A fixed set of generic artifact pitfalls is appended so that severe corruptions are always explicitly penalized.

Rubric Evaluation. Given a sampled prediction $\hat{x}$ and rubric $\mathcal{R}$, the judge evaluates each criterion independently and returns binary pass/fail decisions. Let $\mathcal{P}_i = \{j : p_{ij} = 1, w_j > 0\}$ denote satisfied non-pitfall items and $\mathcal{F}_i = \{j : p_{ij} = 0, w_j < 0\}$ the violated pitfalls. The aggregate reward is:

$$R_i^{\text{rubric}} = \text{clip}\left(\frac{\sum_{j \in \mathcal{P}_i} w_j + \sum_{j \in \mathcal{F}_i} w_j}{\sum_{j: w_j > 0} w_j}, 0, 1\right). \quad (18)$$

This explicit aggregation enables interpretable debugging—we know *which* checklist items fire on each sample—and tends to stabilize GRPO training, at the cost of two judge calls per instance (rubric generation + evaluation). See prompt for RaR generation and evaluation in Figure 17.

D.3. Explicit Error Counting

To directly tie training to concrete visual failure modes, we test an explicit error-count reward. The judge is asked to list all distinct errors in the predicted garment region (plus any global artifacts), returning a set of short labels with explicit severity weights annotations across each evaluation aspect:

$$\mathcal{E}_a^{(i)} = \{(e_k, w_k)\}_{k=1}^{|\mathcal{E}_a^{(i)}|}, \quad w_k \in \{0, 1\}. \quad (19)$$

Weighting follows a structured policy: 1 errors meaningfully break garment transfer or realism (missing parts, garment-type mismatch, large length/placement errors, obvious color/pattern mismatch, identity/pose change, significant leaks/halos/occlusions), while all less detrimental errors are weighted closer to 0.1.

The raw error count and reward are computed as:

$$R_a^{(i)} = \frac{1}{|\mathcal{E}_a^{(i)}|} \sum_k w_k \quad \text{where} \quad R_i = 1 - \frac{1}{5} \sum_{a=1}^5 R_a^{(i)} \quad (20)$$

Raw counts are not directly comparable across prompts, so for each GRPO group, we apply the same group-wise calibration in Section 2.2 to obtain R_i^{explicit} . This representation is more structured but still suffers from per-image subjectivity: near-duplicate generations can receive different error sets due to VLM stochasticity in label selection as shown in the poor results, see Table 4. See prompt for EEC in Figure 17.

D.4. Cascaded Error Counting for Evaluation

While the above methods are employed for GRPO training, our *primary* offline evaluation metric is **cascaded error counting**. Plain error counting exhibits high variance because the label space and severities can drift across images—the VLM may use slightly different phrases or overlook common errors inconsistently. Cascaded error counting reduces this variance by *sharing* an evolving error vocabulary across candidates, building a pool of canonical error labels for each $(x_{\text{src}}, x_{\text{ref}})$ pair.

Pool-Driven Evaluation. We process candidates in blocks of size $B \ll K$ to keep context manageable. For the first block, we run a two-pass cascade:

1. **Discover.** The VLM returns preliminary per-image error sets (with severity labels) and a list of *typical errors* observed across the block. We canonicalize labels (lowercasing, synonym normalization) and merge into an initial pool $\mathcal{P}^{(0)}$ of canonical error labels with severities.
2. **Refine.** We re-evaluate the same images, conditioning on the current pool $\mathcal{P}^{(0)}$ and the previous errors as hints. The VLM is instructed to prefer pool labels when describing the same issue, merge near-duplicates to pool wording, and add missing errors it now notices. This yields refined per-image error sets $\mathcal{E}_i^{(1)}$ and an updated pool $\mathcal{P}^{(1)}$.

Subsequent blocks are evaluated in a single POOL-EVAL pass, conditioned on the current pool, with limited allowance for new labels. New error suggestions are merged into the pool, which continues to evolve across blocks. See Figures 20 and 21 for prompt details

Error Canonicalization. We apply a deterministic normalizer ϕ to all error labels: lowercasing, removing non-alphanumeric characters, collapsing whitespace, and synonym replacement (e.g., “bleed” $\rightarrow$ “leak”, “haloing” $\rightarrow$ “halo”). Duplicates are merged by label, keeping the more severe annotation if severities disagree.

Cyclical Refinement. For comprehensive evaluation, we optionally run a *cyclical* variant: after the first full pass (pool evolves), we execute a second full pass using the *final* pool from pass-1. The second pass runs REFINE on every block with the stabilized vocabulary, then unions errors from both passes (with major severity dominating on conflicts). This two-pass approach further reduces variance while ensuring no errors are missed due to early-pool incompleteness.

Final Score Computation. The cascaded error sets $\mathcal{E}_i$ are converted to counts C_i via Equation (20). Each error in the pool has a severity and weight (major vs. minor), and we summarize each prediction by its weighted error count over the globally consistent error taxonomy. Final scores are weighted counts $C_i = \sum_{(e, \text{sev}) \in \mathcal{E}_i} w(\text{sev})$ with $w(\text{major}) = 5$ and $w(\text{minor}) = 1$.

Alignment with Training Reward. Implicit error counting, our primary training reward, is designed to align with the cascaded error evaluation metric: both rely on fine-grained, judge-generated error analysis rather than coarse holistic impressions, but IEC operates per-image during training while cascaded counting operates across candidates during evaluation. In practice we train GRPO with implicit error counting (and compare against direct Likert, rubric, and explicit error rewards), and report cascaded error as our main quantitative metric for comparing different policies after training. The shared error vocabulary ensures that evaluation scores are directly comparable across model checkpoints and baselines, unlike per-image methods where label drift can confound comparisons.

Prompt for Direct Scalar Reward

System Prompt:

You are an expert fashion-photo inspector evaluating a single virtual try-on (VTO) result. You will see:

- Source_image – the original photo of the person
- Reference_image – the target garment
- Predicted_image A – the VTO output that should place the Reference garment on the Source
- optional Source_mask – the garment region of interest on the Source
- optional Garment_category $\in$ {upper_body/top, lower_body/lower, dress}

Your task is to assign a single scalar reward in [0,1] for Predicted_image A, Judging using the these priorities:

1. **Garment-transfer fidelity / accuracy**
 - correct placement on body; sleeve & hem length; collar/neckline; pattern orientation; size vs pose.
2. **Attribute preservation**
 - colors, dominant pattern, and texture match the Reference garment.
3. **Physical realism & integration quality**
 - natural drape and wrinkles; fabric follows pose; clean boundaries (no halos, leaks, distortion, or blur).
4. **Lighting consistency**
 - garment lighting, shadows, and highlights match the Source scene.
5. **Source integrity**
 - face, hair, and background remain unchanged except where naturally covered by the garment.

POLICY

- If Source_mask is provided, focus primarily on pixels inside that region; still penalize clear global artifacts (identity/pose change, duplicated limbs, background corruption, obvious color bleed).
- Garment_category tells you which region should be edited; penalize if a different primary garment type was modified.
- Be strict but not pedantic: a single tiny nit does not make the image unusable.

Return exactly one JSON object, with this structure and nothing else:

```
{
  "reward": 0.0-1.0,
  "aspect_scores": { "garment_transfer": 0.0-1.0, "attribute_preservation": 0.0-1.0, "realism_integration": 0.0-1.0, "lighting_consistency": 0.0-1.0, "source_integrity": 0.0-1.0 },
  "notes": ["<= 15-word overall summary", "optional short extra note"]
}
```

Guidelines:

- Set aspect_scores first; then let "reward" be roughly their mean, but you may bias it DOWN if any aspect is catastrophically bad (e.g., wrong garment, severe corruption).
- Clamp all scores to [0,1].
- Do NOT include any text outside the JSON object. No code fences.

Figure 16. **Direct Reward Prompt:** Details the system prompt used for generating Direct rewards with each predicted image evaluated individually.

Prompt for Rubrics as a Reward (RaR)

RaR Rubric Generation System Prompt:

You are an expert rubric writer for Virtual Try-On (VTO) judging. Given Source_image (person), Reference_image (garment), and optional masks (garment region or inpainted source), produce a self-contained rubric tailored to this pair.

Rules:

- 18-30 items total.
- Each item: { "id": "R1", "title": "2-4 words", "description": "One sentence, self-contained", "category": "**Essential|Important|Optional|Pitfall**", "weight": number }
- Weights: Essential=5, Important=3-4, Optional=1-2, Pitfall=-2 to -5 (choose severity by real-world impact).
- Cover: transfer fidelity (placement, sleeve/hem length, collar/neckline, proportion), attribute preservation (color/pattern/texture), realism/integration (edges, halos, leaks, drape), lighting consistency, and source integrity (human features, skin color, pose).
- If a mask or inpainted source is present, add 1-2 mask-consistency criteria (no leaks beyond mask, coverage respects mask).
- Mandatory final section - "Global Artifact & Corruption" (8-12 Pitfall items): general, image-agnostic pitfalls to catch bad generations (e.g., background corruption or cloning; floating cutouts/ghosting; texture tiling/repetition; severe warping near joints; hard erasure halos; duplication or truncation of limbs; heavy blur/smear; banding/compression blocks; random brush strokes/noise; occlusion order errors; shadow/hard-cut mismatch; garment pixels staining environment). Phrase positively as "Avoids X ..."; violation triggers penalty.
- Write items so they apply consistently across any predicted image for this pair.
- Output only a JSON array of rubric items. No extra text.

RaR Rubric Evaluation System Prompt:

You are an expert VTO judge. Task: Given Source_image, Reference_image, optional masks, a Predicted_image, and a rubric list, evaluate each criterion independently and compute a normalized reward.

Return exactly:

```
{
  "per_item": [{"id": "R1", "pass": true|false, "note": "≤10 words"}, ...],
  "aggregates": {
    "sum_positive_weights": float,
    "sum_achieved": float,
    "sum_penalties": float,
    "normalized_reward": float // (sum_achieved + sum_penalties) / sum_positive_weights, clipped to [0,1]
  }
}
```

Rules:

- For Essential/Important/Optional: if satisfied, add its weight to sum_achieved; else add 0.
- For Pitfall: if violated, add its (negative) weight to sum_penalties; if avoided, add 0.
- sum_positive_weights = sum of weights over non-Pitfall items only.
- normalized_reward = clip((sum_achieved + sum_penalties) / sum_positive_weights, 0, 1).
- Be strict on "Global Artifact & Corruption" pitfalls: background corruption, ghosting, tiling, severe warping, halos/erasure, duplicated/truncated limbs, heavy smearing/blur, compression artifacts, random strokes/noise, occlusion order errors, environment staining.
- utilize the rubric text literally for evaluations.
- Return only the JSON.

Figure 17. **RaR Prompt:** Details the system prompt used to generate rubrics given the source, reference, mask, and garment conditioning and to evaluate each predicted image individually on the generated rubrics.

Prompt for Explicit Error Counting (EEC) Reward

System Prompt:

You are an expert fashion-photo inspector evaluating a single virtual try-on (VTO) result. You will see:

- Source_image – the original photo of the person
- Reference_image – the target garment
- Predicted_image A – the VTO output that should place the Reference garment on the Source
- optional Source_mask – the garment region of interest on the Source
- optional Garment_category $\in$ {upper_body/top, lower_body/lower, dress}

Your task is to **identify and count errors by aspect and weight each error across** these priorities:

1. **Garment-transfer fidelity / accuracy**
 - correct placement on body; sleeve & hem length; collar/neckline; pattern orientation; size vs pose.
2. **Attribute preservation**
 - colors, dominant pattern, and texture match the Reference garment.
3. **Physical realism & integration quality**
 - natural drape and wrinkles; fabric follows pose; clean boundaries (no halos, leaks, distortion, or blur).
4. **Lighting consistency**
 - garment lighting, shadows, and highlights match the Source scene.
5. **Source integrity**
 - face, hair, and background remain unchanged except where naturally covered by the garment.

POLICY

- If Source_mask is provided, focus primarily on pixels inside that region; still penalise clear global artifacts (identity/pose change, duplicated limbs, background corruption, obvious colour bleed).
- Garment_category tells you which region should be edited; penalise if a different primary garment type was modified.
- Be strict but not pedantic: a single tiny nit does not make the image unusable.
- For error, set "weight": [0-1]. weight close to 1 for errors that meaningfully break garment transfer or realism (missing parts; clear garment-type mismatch; large length/placement error; obvious color/pattern mismatch; identity/pose change; big leaks/halos/occlusions), otherwise weight closer to 0 for very minor issues and inbetween that spectrum based on the severity of error.

Return exactly one JSON object, with this structure and nothing else:

```
{
  "errors": [
    {
      "label": "short error description",
      "aspect": "garment_transfer|attribute_preservation|realism_integration|lighting_consistency|source_integrity",
      "weight": "0.0-1.0",
    }, ... ]
  }
```

Guidelines:

- Do NOT include any text outside the JSON object. No code fences.

Figure 18. **Explicit Error Counting Reward Prompt:** Details the system prompt used for counting and weighing errors on each predicted image.

Prompt for Implicit Error Counting (IEC) Reward

System Prompt:

You are an expert fashion-photo inspector evaluating a single virtual try-on (VTO) result. You will see:

- Source_image – the original photo of the person
- Reference_image – the target garment
- Predicted_image A – the VTO output that should place the Reference garment on the Source
- optional Source_mask – the garment region of interest on the Source
- optional Garment_category $\in$ {upper_body/top, lower_body/lower, dress}

Your task is to **identify and count errors by aspect and severity, but DO NOT return counts.**

Instead, convert them into per-aspect rewards in [0,1], then a final_reward in [0,1] for Predicted_image A, where:

- 1.0 = almost perfect try-on, no meaningful errors
- 0.8–1.0 = strong result with only minor error
- 0.5–0.79 = mixed quality; clear errors but still usable
- 0.2–0.49 = poor; serious errors in garment transfer or realism
- 0.0–0.19 = failed try-on; garment identity wrong or image badly corrupted

Judge using the these priorities:

- Garment-transfer fidelity / accuracy**
 - correct placement on body; sleeve & hem length; collar/neckline; pattern orientation; size vs pose.
- Attribute preservation**
 - colors, dominant pattern, and texture match the Reference garment.
- Physical realism & integration quality**
 - natural drape and wrinkles; fabric follows pose; clean boundaries (no halos, leaks, distortion, or blur).
- Lighting consistency**
 - garment lighting, shadows, and highlights match the Source scene.
- Source integrity**
 - face, hair, and background remain unchanged except where naturally covered by the garment.

POLICY

- If Source_mask is provided, focus primarily on pixels inside that region; still penalise clear global artifacts (identity/pose change, duplicated limbs, background corruption, obvious colour bleed).
- Garment_category tells you which region should be edited; penalise if a different primary garment type was modified.
- Be strict but not pedantic: a single tiny nit does not make the image unusable.

Return exactly one JSON object, with this structure and nothing else:

```
{
  "reward": 0.0-1.0,
  "aspect_rewards": { "garment_transfer": 0.0-1.0, "attribute_preservation": 0.0-1.0, "realism_integration": 0.0-1.0, "lighting_consistency": 0.0-1.0, "source_integrity": 0.0-1.0 },
  "errors": [
    {
      "label": "short error description",
      "aspect": "garment_transfer|attribute_preservation|realism_integration|lighting_consistency|source_integrity"
    }, ... ]
}
```

Guidelines:

- Set aspect_scores first; then let "reward" be roughly their mean, but you may bias it DOWN if any error/aspect is catastrophically bad (e.g., wrong garment, severe corruption).
- Do NOT include any text outside the JSON object. No code fences.

Figure 19. **Implicit Error Counting Reward Prompt:** Details the system prompt used for implicitly counting errors and determining aspect scores on each predicted image.

Prompt for Cascaded Error Counting (CEC) Metric

System Prompt:

You are an expert fashion-photo inspector for Virtual Try-On (VTO). You will receive:

- Source_image (person)
- Reference_image (garment)
- N Predicted_images labeled A, B, C, ...
- Optional Source_mask (focus the garment region)
- Garment_category $\in$ {upper_body/top, lower_body/lower, dress}

POLICY

- 1) Judge inside the Source_mask. Outside-of-mask judge obvious global artifacts
- 2) Compare logically: The Predicted garment must come from Reference and be correctly applied to the Source body/pose. Do not compare prediction to source or prediction to reference in isolation.
- 3) Granularity: List distinct visual errors as short labels (≤ 8 words each).
- 4) Severity: For every error, set "severity": "major" | "minor". Use major for errors that meaningfully break garment transfer or realism (missing parts; clear garment-type mismatch; large length/placement error; obvious color/pattern mismatch; identity/pose change; big leaks/halos/occlusions). Otherwise use minor.
- 5) Count rule (the "count" you return must match these weights):
 - "major" contributes $\llcorner \text{MAJOR_WEIGHT} \gg$
 - "minor" contributes +1

SCHEMA (JSON only; no extra text)

• DISCOVER / REFINE:

```
{
  "candidates": {
    "A": { "errors": [{"label": "...", "severity": "major|minor"}, ...], "count": int},
    "B": { "errors": [...], "count": int},
    ...
  },
  "typical_errors": [{"label": "...", "severity": "major|minor"}, ...],
  "schema": "cascaded_error_v2"
}

• POOL-EVAL (when a Known_error_pool is provided):
{
  "candidates": {
    "A": { "errors": [{"label": "...", "severity": "major|minor"}, ...], "count": int},
    "B": { "errors": [...], "count": int},
    ...
  },
  "new_error_suggestions": [{"label": "...", "severity": "major|minor"}, ...],
  "schema": "cascaded_error_pool_v2"
}
```

- Focus on the transferred garment in the predicted image from the reference.
- Check the predicted image for the correct source image pose, identity and non-mask area.
- Be as detailed and fine-grained as possible e.g what hem is incorrect and how, what artifacts are made to the source garments that ought not to be etc
- Look for where the imagination of the reference garment on the source breaks human expectations
- If no source_mask is provided, infer the garment region from garment_category and the visible garment pixels in the predicted_image, and judge within that inferred region
- Count clear global artifacts.

Figure 20. **Cascaded Error Counting Reward Prompt:** Details the system prompt used for counting and obtaining severity classification for each error.

User Prompts for Cascaded Error Counting (CEC) Metric

Pool Phase User Prompt:

Task: For each candidate, list errors (short labels) and compute 'count' using the policy weights.
Return JSON with "candidates" and "typical_errors" (de-duplicated, sorted by prevalence).
Garment_category: {garment_category}

Verify Phase User Prompt:

You will receive:

- previous_errors (per candidate): your earlier labels to correct/update
- Known_error_pool: canonical labels seen so far

Tasks:

- Re-check each candidate carefully.
 - Add missing errors ones you now notice, if any.
 - Prefer pool labels when they describe the same issue; merge near-duplicates to the pool wording.
 - Assign severity (major/minor) per error. Recompute 'count' with the policy weights.
- Return JSON with "candidates" and "typical_errors" (de-duplicated, sorted by prevalence). Keep labels short and reusable.
Garment_category: {garment_category}

Known_error_pool (JSON): {known_pool_json}

previous_errors (JSON): {previous_errors_json}

POOL-EVAL Phase User Prompt:

You are given a Known_error_pool of canonical error labels observed on typical VTO predicted images for this src and ref pair.

For each candidate:

- Reuse labels from the pool when applicable.
- You may add ****new**** labels if strictly needed; keep them short and reusable.
- Every error must include the keys "label" and "severity" (severity must be "major" or "minor").
- Ensure "count" follows the severity weights described in the system prompt.

Return exactly this JSON (no extra text):

```
{
  "candidates": {
    "A": {
      "errors": [
        {
          "label": "...",
          "severity": "major|minor",
          "count": 0
        },
        {
          "label": "...",
          "severity": "major|minor",
          "count": 0
        }
      ],
      "new_error_suggestions": [
        {
          "label": "...",
          "severity": "major|minor"
        }
      ],
      "schema": "cascaded_error_pool_v2"
    }
  }
}
```

Garment_category: {garment_category}

Known_error_pool (JSON): {known_pool_json}

Figure 21. Cascaded Error Counting Reward Prompt: Details the user prompt used for each phase of CEC