

# *CellFluxRL*: Biologically-Constrained Virtual Cell Modeling via Reinforcement Learning

Dongxia Wu<sup>1\*</sup>, Shiye Su<sup>1\*</sup>, Yuhui Zhang<sup>1\*</sup>, Elaine Sui<sup>1</sup>, Emma Lundberg<sup>1</sup>, Emily B. Fox<sup>1†</sup>, and Serena Yeung-Levy<sup>1†</sup>

<sup>1</sup>Stanford University, Stanford, CA, USA


{dowu, shiye, yuhui, esui, emmalu, ebfox, syyeung}@stanford.edu

**Abstract.** Building virtual cells with generative models to simulate cellular behavior *in silico* is emerging as a promising paradigm for accelerating drug discovery. However, prior image-based generative approaches can produce implausible cell images that violate basic physical and biological constraints. To address this, we propose to post-train virtual cell models with reinforcement learning (RL), leveraging biologically meaningful evaluators as reward functions. We design seven rewards spanning three categories—biological function, structural validity, and morphological correctness—and optimize the state-of-the-art *CellFlux* model to yield *CellFluxRL*. *CellFluxRL* consistently improves over *CellFlux* across all rewards, with further performance boosts from test-time scaling. Overall, our results present a virtual cell modeling framework that enforces physically-based constraints through RL, advancing beyond “visually realistic” generations towards “biologically meaningful” ones.

## 1 Introduction


Wet-lab experiments have long been the primary bottleneck in drug discovery. To conduct a single experiment, researchers must not only purchase expensive reagents, consumables, and equipment, but also wait weeks or even months to synthesize drugs and culture cells. Consequently, biologists have long envisioned building virtual cells [4, 16, 36] to accelerate this process by simulating cell behavior *in silico*. With recent advances in generative modeling, such as diffusion models [13, 38, 39] and flow matching [22, 23, 26, 27], alongside high-throughput image-based screening technologies [7, 8, 12] that produce terabytes of data, there is growing interest in developing image-based virtual cells to predict morphological responses to perturbations.

Despite its success, we observe that these image-based virtual cell models can produce images that look realistic yet are biologically implausible. For instance, using the state-of-the-art image-based virtual cell model, *CellFlux* [46], we observe anomalies such as the cell nucleus being generated outside of the cytoplasm (Figure 1). Such violations greatly limit real-world use of virtual cell models, reducing their practical value.



**Fig. 1:** Failure of cell generation.

\* Equal contribution. † Equal advising.



**Fig. 2: Motivation.** Current generative models for simulating cellular perturbations can fail to produce physically plausible cell images. For example, nuclei may appear outside the cell membrane. We design a suite of biologically meaningful verifiers in three roles: (1) as *evaluators* to assess the biological correctness of generated images, (2) as *reward signals* to improve generation via reinforcement learning, and (3) as *verification modules* to enhance sample quality through test-time scaling.

We hypothesize that these failures arise from a mismatch between the pixel-level flow-matching objective and global physical constraints. Flow-based generative models [23] learn velocity fields that transport samples from one distribution to another. While this objective can encourage visual fidelity, it does not explicitly enforce higher-level constraints on cell structure and properties. This may yield samples that appear plausible locally but are incorrect globally.

To address this issue, we propose *CellFluxRL*, applying reinforcement learning (RL)-based post-training with biologically-meaningful evaluators as reward functions. Concretely, many constraints of interest admit existing or readily obtained evaluators that can efficiently assign scores to generated images (e.g., whether a sample satisfies structural or functional criteria). Although these evaluators are typically non-differentiable and thus cannot be optimized via standard backpropagation, they can be used as reward signals for RL [24, 47]. At a high level, our RL procedure alternates between: (1) **sampling**, where we generate multiple images and evaluate them using biologically-meaningful evaluators; and (2) **optimization**, where we increase the likelihood of high-reward samples and decrease the likelihood of low-reward samples. This directly penalizes physically invalid generations while reinforcing physically consistent ones. Although our work focuses on improving *CellFlux*, a flow-based model, the same training ideas can be applied to other generative modeling frameworks.

To demonstrate the potential of this approach, we post-train *CellFlux* to yield *CellFluxRL*, optimizing seven reward models that assess correctness across three key categories: **biological function**, such as mode of action; **structural properties**, including nuclear roundness and containment of the nucleus within the cytoplasm; and **morphological properties**, like the size and number of nuclei and the size of the cytoplasm. Our overall reward is a weighted linear combination of component rewards. Evaluating the learned model across all reward components, *CellFluxRL* achieves higher scores than the base model. These results indicate that RL provides an effective mechanism for aligning image-based generative models with biologically-meaningful constraints. We additionally present



qualitative case studies and ablations to analyze how RL hyperparameters and reward design influence performance.

Moreover, these reward models enable a simple form of test-time scaling. Following the selection strategy of Ma *et al.* [28], we sample  $N$  candidate images for a given condition and select the one with the highest reward. We observe monotonic improvements as  $N$  increases, indicating a clear test-time scaling trend, and achieving an additional gain over single-sample generation.

Our work also shines a light on limitations of existing evaluation metrics for image-based generative models of cellular morphology. To date, focus has been on image generation quality scores, like FID and KID, rather than biological validity. Our considered rewards not only enable RL-based training, but also new, biologically meaningful evaluation criteria.

In summary, our work identifies a key limitation of existing image-based virtual cell modeling: the lack of explicit enforcement of physical correctness under pixel-level training objectives. We address this limitation by incorporating biologically-meaningful evaluators through RL, which substantially improves physical plausibility of generated images. We further show that these gains can be amplified via test-time scaling. Finally, our rewards serve as new metrics for the community to benchmark against. Overall, our results move virtual cell image generation from “looking good” to being *biologically consistent*, supporting the downstream goal of drug discovery and personalized medicine with reduced reliance on costly wet-lab experiments.

## 2 Related Work

*Virtual cell modeling.* A *virtual cell* model aims to simulate cellular responses to perturbations *in silico*. This can be formulated as a generative modeling problem  $p(x_1|x_0, c)$ , where  $x_0$  and  $x_1$  denote cell states before and after perturbation  $c$ . Early virtual cell models relied on mathematical and physics-based frameworks, such as systems of ordinary differential equations (ODEs) [35, 40, 45]. However, their expressiveness is inherently limited by the strict assumptions and simplified dynamics that these formulations require. With recent advances in generative modeling and image-based high-throughput screening, a growing body of work has explored using VAEs, GANs, diffusion models, and flow matching to simulate cellular state transitions [1, 3, 15, 18, 32, 34, 46]. A notable paradigm is *CellFlux* [46], which reformulates the problem as a distribution-to-distribution transformation and leverages flow matching to mitigate batch effects, achieving state-of-the-art performance. Nevertheless, we observe that while these models generate visually plausible cells, the generations can violate fundamental biological constraints. To address this challenge, we introduce RL post-training guided by biologically meaningful rewards, encouraging virtual cell models to generate more physically and biologically plausible cell images.

*Physics-aware generative models.* The growing interest in utilizing generative models as “world simulators” or “world models” necessitates strict adherence to

physical accuracy [9, 29, 41]. However, recent studies indicate that state-of-the-art image and video generative models still struggle to maintain physical correctness, a limitation that cannot be resolved simply by scaling up data and model parameters [17, 30]. To tackle this challenge, the community has proposed various approaches: modifying architectures to inject physical constraints [14, 44], altering the inference process by using LLMs for generation planning [21, 48], or modifying training objectives to incorporate physical preferences [5, 25, 33, 42]. Our approach aligns with objective-based methods; in virtual cell modeling, we can readily define numerous biologically meaningful constraints. Optimizing these constraints requires minimal architectural modifications while yielding robust results. Furthermore, these rewards facilitate test-time scaling to continuously enhance generation quality.

*Reinforcement learning for generative modeling.* While generative models based on diffusion or flow matching achieve high visual fidelity, aligning them with strict physical constraints remains an open challenge. RL has emerged as a powerful solution [2, 11, 19, 24, 43], yet its development is still in its nascent stages. The primary hurdle lies in evaluating the exact likelihood of samples, necessary for adjusting generation probabilities based on rewards, which is generally intractable in diffusion and flow matching frameworks. Prior works, such as FlowGRPO [24] and DanceGRPO [43], formulate this as a Markov Decision Process (MDP) to derive probabilities, whereas recent algorithms like DiffusionNFT [47] introduce forward processes to estimate these likelihoods. Our work builds upon the DiffusionNFT-style forward probability estimation process, which offers rapid estimation and state-of-the-art results. Crucially, it accommodates our distribution-to-distribution flow matching (from original to perturbed states), whereas MDP-based approaches assume a standard Gaussian prior.

### 3 Cellular Perturbation Modeling

Our approach takes as input a pretrained generative model for cellular morphology prediction and post-trains it using RL to improve physical correctness. In this section, we describe the perturbation modeling problem and the properties required for the base model.

*Objective.* Let  $\mathcal{X} \subseteq \mathbb{R}^{H \times W \times C}$  denote the space of multi-channel fluorescence microscopy images, where each channel highlights a distinct cellular structure (e.g., nucleus, cytoskeleton, mitochondria). Let  $\mathcal{C}$  denote the space of perturbations, encompassing chemical compounds (drugs). Given an unperturbed cell image  $x_0$  and a perturbation  $c \in \mathcal{C}$ , the goal is to learn a generative model that samples from the conditional distribution  $p(x_1 | x_0, c)$ , where  $x_1$  represents the cell’s morphology after treatment. Such a model enables *in silico* simulation of cellular responses that would otherwise require costly wet-lab experiments.

*Data and batch effects.* Cell morphology data are collected via high-content microscopy screening, in which multi-well plates are prepared with both *control* wells (untreated) and *perturbed* wells (treated with a chemical compound). Experiments are conducted across multiple batches, each introducing systematic technical variations—differences in staining intensity, illumination, or imaging conditions—that are unrelated to the perturbation itself. Because imaging is destructive (cells are fixed and stained), paired before-and-after observations of the same cell are unavailable. Instead, we observe unpaired sets of control images  $\{x_0\}$  and perturbed images  $\{x_1\}$  within each batch.

*Requirements for the base model.* These data characteristics impose two key requirements on the base generative model. First, because observations are unpaired, the model must learn a *distributional* transformation from control to perturbed cells rather than a pointwise mapping. Second, because batch effects confound perturbation signals, the model should condition on same-batch control images as its source distribution, so that it learns only the perturbation-induced changes rather than technical artifacts. While our approach is compatible with any such base model, we instantiate it as *CellFlux* [46]. *CellFlux*, based on flow matching, offers a natural framework satisfying the above requirements: it learns a velocity field  $v_\theta: \mathcal{X} \times [0, 1] \times \mathcal{C} \rightarrow \mathcal{X}$  that continuously transports the control distribution  $p_0$  to the perturbed distribution  $p_1(\cdot | c)$  within each batch. Training proceeds by sampling pairs  $(x_0, x_1)$  from source and target distributions, constructing linear interpolations  $x_t = (1 - t)x_0 + tx_1$  for  $t \sim \mathcal{U}[0, 1]$ , and minimizing:

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{x_0 \sim p_0, x_1 \sim p_1, t \sim \mathcal{U}[0, 1]} \|v_\theta(x_t, t, c) - (x_1 - x_0)\|_2^2. \quad (1)$$

At inference, a control image  $x_0$  is transformed into a predicted perturbed image  $\hat{x}_1$  by solving the ODE  $dx_t = v_\theta(x_t, t, c) dt$  from  $t=0$  to  $t=1$ . We denote the resulting pretrained model as  $v_\theta$ , which serves as the input to our method.

## 4 Method

In this section, we present our approach to improving the physical and biological fidelity of virtual cell models. We first introduce the formal objective for reward-driven generation (§4.1). Next, we detail the suite of reward models/evaluators (§4.2) designed to capture biological and structural properties. We then describe the RL algorithm (§4.3) used to post-train the base model via contrastive updates. Finally, we discuss test-time scaling (§4.4) to further enhance generation quality at inference time.

### 4.1 Reward-Driven Generation

While the flow matching objective trains the base model to produce images that are distributionally close to real perturbed cells, it operates at the pixel-level

and does not explicitly encourage higher-level physical or biological correctness. In practice, this can lead to generated images that appear visually plausible but violate known cellular properties—for example, producing nuclei of incorrect size, cells with implausible morphology, or images that fail to reflect the expected mode of action (MoA) of a compound (Figure 2). Such violations limit the utility of virtual cell models for downstream applications like drug screening, where physical correctness is essential.

We observe that many properties of interest can be assessed by *evaluator functions* that score generated images along specific axes of physical or biological validity. Formally, let  $\{r_k\}_{k=1}^K$  denote a set of reward functions, where each  $r_k: \mathcal{X} \times \mathcal{C} \rightarrow \mathbb{R}$  measures how well a generated image  $\hat{x}_1$  satisfies a particular property given perturbation  $c$ . These rewards capture complementary aspects of cell structure and function, such as whether the nuclear size is consistent with the expected effect of the perturbation, or whether the overall morphological profile matches the correct MoA. We describe the specific reward functions used in this work in §4.2.

Given the pretrained base model  $v_\theta$ , our objective is to obtain a post-trained model  $v_{\theta'}$  that generates images with higher physical fidelity as measured by  $\{r_k\}$ , while remaining close to the pretrained model. Concretely, we seek to solve:

$$\max_{\theta'} \mathbb{E}_{x_0 \sim p_0, c \sim p(c)} \left[ \sum_{k=1}^K w_k r_k(\hat{x}_1, c) \right] \quad \text{s.t.} \quad D_{\text{KL}}(v_{\theta'} \| v_\theta) \leq \epsilon, \quad (2)$$

where  $\hat{x}_1$  is the image generated by  $v_{\theta'}$  from input  $(x_0, c)$ ,  $w_k$  is the weight of each reward  $r_k(\cdot)$ , and  $D_{\text{KL}}(\cdot \| \cdot)$  is the KL divergence between the post-trained and pretrained models. This regularization is important both to preserve the visual quality and diversity learned during pretraining, and to mitigate *reward hacking*—generating images that achieve high scores through degenerate solutions rather than genuine physical correctness.

Critically, optimizing these rewards is complementary to, and not in conflict with, the original distributional objective. The reward functions encode properties that the *ground truth* perturbed distribution satisfies. For instance, real cells treated with a microtubule destabilizer exhibit smaller, fragmented nuclei, and real cells have mode-of-action-consistent morphological profiles. By directly optimizing for these properties, RL can improve the match between generated and real distributions along biologically meaningful axes that the pixel-level flow matching loss may underweight. In this sense, the rewards provide a form of *targeted distributional alignment*, focusing model capacity on the aspects of the distribution that matter most for physical correctness.

## 4.2 Reward Models

We design a suite of reward functions that capture complementary aspects of biological validity in generated cell images. These rewards fall into three categories: *biological function* rewards that assess whether generated images reflect

the expected biological effects of a perturbation, *structural* rewards that enforce known physical relationships between cellular components, and *morphological* rewards that encourage generated images to match the size and density statistics of real cells. Together, these rewards provide multi-scale supervision—from high-level biological semantics down to low-level geometric properties—that the pixel-level flow matching objective does not explicitly optimize for.

**Biological Function.** Biological function rewards evaluate whether generated images faithfully reflect the biological effects of the applied perturbation.

*Mode of action (MoA).* A drug’s MoA describes the cellular process it targets, such as microtubule destabilization, DNA replication inhibition, or actin disruption. Because different MoAs produce distinct morphological signatures, the ability of a generated image to be correctly classified by MoA serves as a strong indicator of biological fidelity. We leverage an MoA classifier pretrained on real perturbed images [46] and define the reward as the predicted probability of the ground truth MoA class:

$$r_{\text{MoA}}(\hat{x}_1, c) = p_{\text{cls}}(y_c \mid \hat{x}_1), \quad (3)$$

where  $y_c$  is the ground truth MoA label associated with perturbation  $c$  and  $p_{\text{cls}}$  is the pretrained MoA classifier. Higher values indicate that the generated image exhibits a morphological profile characteristic of the ground truth MoA, making different drug classes more faithfully distinguishable.

**Structural Constraints.** Structural constraint rewards enforce known physical relationships between cellular components.

*Nucleus-in-cytoplasm.* In correctly imaged cells, nuclei are fully enclosed within the cytoplasm. Generated images that violate this containment are physically implausible. We segment nuclei and cytoplasm from the generated images using Cellpose [31] and define the reward as the degree to which all detected nuclei are fully enclosed by cytoplasm. This reward penalizes the failure mode of generative models wherein channel-level structures become spatially inconsistent:

$$r_{\text{Nuc-in-Cyto}}(\hat{x}_1, c) = \frac{\text{area}(\text{nucleus mask} \cap \text{cytoplasm mask})}{\text{area}(\text{cytoplasm mask})}. \quad (4)$$

*Nuclear roundness.* Nuclear shape is a biologically informative feature that varies systematically across MoA classes. For instance, microtubule destabilizers produce fragmented, irregular nuclei, while other perturbations preserve smooth, round nuclear morphology. We compute the roundness reward as:

$$r_{\text{Roundness}}(\hat{x}_1, c) = - \left[ \frac{1}{N_{\text{Nu}}} \sum_{i=1}^{N_{\text{Nu}}} \frac{4\pi \cdot \text{area}_i}{\text{perimeter}_i^2} - \mu^{(y_c)} \right]^2 / [\sigma^{(y_c)}]^2, \quad (5)$$

where  $N_{\text{Nu}}$  is the number of identified nuclei, and  $\mu^{(y_c)}$  and  $\sigma^{(y_c)}$  are the average and standard deviation roundness of the MoA-conditioned ground truth distribution of  $\hat{x}_1$  with label  $y_c$ . This encourages the model to produce nuclear shapes consistent with the known morphological effects of each perturbation category.

**Morphological Statistics.** These rewards ensure that the size and density of cellular components in generated images match those of the ground truth distribution, conditioned on the MoA. Unlike the structural rewards above, which enforce qualitative relationships, these rewards target quantitative statistics of the cell population.

For each of the following statistics  $s$ —maximum nucleus size, maximum cytoplasm size, nucleus count, and cytoplasm count—we compute the value  $s(\hat{x}_1)$  from the generated image and compare it to the ground truth distribution for the corresponding MoA class. Let  $\mu_s^{(y_c)}$  and  $\sigma_s^{(y_c)}$  denote the mean and standard deviation of statistic  $s$  across real images with MoA label  $y_c$ . We define the reward as the negative normalized deviation:

$$r_s(\hat{x}_1, c) = -\frac{[s(\hat{x}_1) - \mu_s^{(y_c)}]^2}{[\sigma_s^{(y_c)}]^2}, s \in \{\text{NucSize}, \text{CytoSize}, \text{NucCount}, \text{CytoCount}\}. \quad (6)$$

This formulation penalizes generated images whose statistics deviate from the ground truth mean, scaled by the natural variability within each MoA class. The four statistics capture complementary aspects of cell morphology:


- **Maximum nucleus size** and **maximum cytoplasm size** reflect the largest cellular components in the image, capturing whether the model produces cells of appropriate scale for the given perturbation. We use the maximum rather than the average to mitigate the potential issue of partially visible cells biasing the statistics.
- **Nucleus count** and **cytoplasm count** measure cell density, ensuring the model does not hallucinate or omit cells relative to what is observed under each treatment condition.

Together, these four rewards encourage the generated cell population to be statistically consistent with real observations in terms of both individual cell size and overall cell density.

### 4.3 Reinforcement Learning on Biological Rewards

We now describe how we optimize the pretrained base model  $v_\theta$  with respect to the reward functions introduced in §4.2.

*Algorithm overview.* Our objective is to maximize the combined reward defined in Eq.(2) (§4.1). Since these biological reward functions are non-differentiable, standard backpropagation is inapplicable, necessitating an RL approach. The



**Fig. 3: CellFluxRL algorithm.** RL post-training seeks to increase the likelihood of high-reward samples and decrease the likelihood of low-reward samples. Therefore, the core training loop of CellFluxRL consists of interleaved phases of sampling and training. (a) Sampling: we generate multiple rollouts from a fixed control image and perturbation condition, scoring each with the reward models and aggregating them into a combined reward  $r$  (Eq. 10), which is group-normalized into an optimality reward  $r_{\text{adv}}$  (Eq. 11). (b) Training: because exact likelihoods in flow matching are intractable, we construct positive and negative velocities from the batch of rollouts and optimize them contrastively to achieve this goal, following DiffusionNFT [47].

core principle of our RL strategy is to increase the generation likelihood of high-reward samples while penalizing low-reward ones. We adopt DiffusionNFT [47], a state-of-the-art online RL algorithm for flow matching. It operates on the flow’s forward process, avoiding intractable log likelihoods, and is constructed from *distribution-agnostic* components, hence extending naturally to our source-to-target flow matching setting without modification.

*Algorithm detail.* At each iteration, DiffusionNFT collects a batch of generated images, evaluates them with respect to the reward functions, and uses the rewards to define an improvement direction over the current policy. The key idea is to split generated samples into *positive* (*high-reward*) and *negative* (*low-reward*) subsets and learn a contrastive update that moves the model towards the positive distribution. Concretely, given a generated image  $\hat{x}_1$  with optimality reward  $r_{\text{adv}} \in [0, 1]$ —a group-normalized score (Eq. 11) that measures how close a sample is to being optimal within its rollout group, and that determines its weight as a positive versus negative example—the training objective is:

$$\mathcal{L}(\theta) = \mathbb{E}_{c, \pi^{\text{old}}(\hat{x}_1|c, x_0), t} \left[ r_{\text{adv}} \|v_\theta^+(x_t, c, t) - v\|_2^2 + (1 - r_{\text{adv}}) \|v_\theta^-(x_t, c, t) - v\|_2^2 \right] + \beta D_{\text{KL}}(v_\theta \| v^{\text{old}}), \quad (7)$$

where  $\beta$  is the KL divergence weight,  $x_t = \alpha_t x_0 + \gamma_t x_1$  is the forward-noised version of the generated image,  $v = \dot{\alpha}_t x_0 + \dot{\gamma}_t x_1$  is the corresponding velocity target, and  $v_\theta^+, v_\theta^-$  are *implicit* positive and negative policies defined as:

$$v_\theta^+(x_t, c, t) := (1 - \gamma) v^{\text{old}}(x_t, c, t) + \gamma v_\theta(x_t, c, t), \quad (8)$$

$$v_\theta^-(x_t, c, t) := (1 + \gamma) v^{\text{old}}(x_t, c, t) - \gamma v_\theta(x_t, c, t). \quad (9)$$

Here  $v^{\text{old}}$  is the data-collection policy (a lagging copy of  $v_\theta$ ), and  $\gamma > 0$  is a hyperparameter controlling the guidance strength. The implicit parameterization



is central to the algorithm: rather than training separate positive and negative models, a single policy  $v_\theta$  is optimized such that its mixture with  $v^{\text{old}}$  simultaneously fits high-reward samples (via  $v_\theta^+$ ) and avoids low-reward ones (via  $v_\theta^-$ ). The optimal solution satisfies  $v_{\theta^*} = v^{\text{old}} + \frac{2}{\gamma}\Delta$ , where  $\Delta$  is the reinforcement guidance direction pointing from the negative towards the positive distribution.

This formulation naturally regularizes the post-trained model towards the pretrained policy: when  $\gamma$  is large, the guidance strength  $\frac{2}{\gamma}$  is small and the model stays close to  $v^{\text{old}}$ ; when  $\gamma$  is small, the model is allowed to deviate more aggressively. The data-collection policy  $v^{\text{old}}$  is itself updated via an exponential moving average of  $v_\theta$ .

*Rollout and advantage estimation.* During sampling, we fix a perturbation condition  $c$  and a source control image  $x_0$ , and generate a group of  $m$  candidate images  $\{\hat{x}_1^{(i)}\}_{i=1}^m$ . Since the base model,  $v_\theta$ , augments  $x_0$  with Gaussian noise before transporting it via the learned velocity field, diversity within each group arises primarily from this stochastic noise injection as well as any ODE discretization error. We find this yields sufficient variation for DiffusionNFT to distinguish positive from negative generations. Each candidate is scored by the  $K$  reward functions, which are aggregated via a weighted sum into a per-image combined reward

$$r(\hat{x}_1, c) = \sum_{k=1}^K w_k r_k(\hat{x}_1, c). \quad (10)$$

The raw combined rewards are then normalized within the group and mapped to optimality rewards

$$r_{\text{adv}}^{(i)} = 0.5 + 0.5 \text{clip}\left(\frac{r^{(i)} - \frac{1}{m} \sum_{j=1}^m r^{(j)}}{Z_c}, -1, 1\right) \in [0, 1], \quad (11)$$

following the advantage normalization scheme [47], where  $Z_c > 0$  is a normalization constant. The forward process is then applied to each generated image, and the loss in Eq.(7) is computed over the group.

#### 4.4 Test-Time Scaling

A key benefit of having explicit reward functions is that they can not only be used for training, but also to select among candidate generations at inference time. This strategy has proven highly effective in reasoning models, where best-of- $N$  selection with a reward model or verifier yields consistent improvements that scale predictably with  $N$  [10, 20, 37]. We apply the same principle to virtual cell generation.

Given a perturbation condition  $c$  and a source control image  $x_0$ , we generate  $N$  candidate images  $\{\hat{x}_1^{(i)}\}_{i=1}^N$  and select the one with the highest reward:

$$\hat{x}_1^* = \arg \max_{i \in \{1, \dots, N\}} r(\hat{x}_1^{(i)}, c). \quad (12)$$

This provides a simple, training-free mechanism to improve generation quality given additional inference compute budget. Moreover, best-of- $N$  selection is complementary to reinforcement learning: RL post-training improves the *base distribution* from which candidates are drawn, so that even modest values of  $N$  yield high-quality outputs, while test-time selection contributes additional gains.

## 5 Results

In this section, we evaluate *CellFluxRL*’s ability to generate biologically faithful cellular images. After detailing our experimental setup (§5.1), we present the main results, demonstrating that RL with biologically-aligned rewards systematically improves performance across all biological metrics (§5.2). We show that test-time scaling enhances performance (§5.3) and provide ablation studies.

### 5.1 Experimental details

*Datasets.* We follow the dataset configuration of the base *CellFlux* model and use the high-content microscopy perturbation dataset BBBC021 [6]. BBBC021 is a chemical perturbation dataset comprising 98K three-channel images at  $96 \times 96$  resolution, collected under 26 chemical perturbations grouped into 12 modes of action (MoA). We follow the train/test split used in CellFlux [46].


*Baselines.* We compare our proposed method against the pretrained base model, *CellFlux* [46], as well as two prior generative baselines, *PhenDiff* [3] and *IMPA* [32].

*Evaluation metrics.* In addition to the rewards reported in §4.2 (higher is better), which measure the biological meaningfulness of the generated cells, we also report standard image quality metrics (FID and KID). These metrics measure the image distribution similarity via Fréchet and kernel-based distances between generated images and ground truth images (lower is better).

*Training details.* We instantiate our method using a pretrained *CellFlux* flow matching model. *CellFluxRL* is post-trained using the online RL algorithm DiffusionNFT [47] to optimize a weighted sum of 7 biologically-constrained rewards:  $r = 5.0 r_{\text{MoA}} + 2.0 r_{\text{Nuc-in-Cyto}} + r_{\text{Roundness}} + r_{\text{NucSize}} + r_{\text{CytoSize}} + r_{\text{NucCount}} + r_{\text{CytoCount}}$ . MoA and Nuc-in-Cyto are empirically harder to optimize, motivating their higher weights. As an alternative that requires no prior knowledge for weight tuning, we also tried a uniform weight of 1.0 for all normalized rewards, which likewise improves over the base model. The training objective includes a KL divergence weight ( $\beta = 1$ ) to regularize the post-trained model against the pretrained policy and prevent reward hacking. All other hyperparameter settings follow DiffusionNFT [47]. At inference, we further enhance performance using test-time scaling (TTS) with a best-of- $N$  selection strategy ( $N = 4$ ), selecting the generated sample with the highest overall reward. Training is conducted for 1200 steps on 1 H100 GPU for 32 hours.

**Table 1:** Quantitative comparison across biological rewards and generative quality metrics. Each row represents a different evaluation metric, and each column represents a different method. TTS is *CellFluxRL* with test time scaling by best-of- $N$  with  $N = 4$ , where the best sample is selected by the weighted total reward. **Bold** values indicate the best performance; underlined values indicate the second best.


| Metric                    | PhenDiff [3] | IMPA [32] | CellFlux [46] | CellFluxRL   | +TTS         |
|---------------------------|--------------|-----------|---------------|--------------|--------------|
| <i>Biological Rewards</i> |              |           |               |              |              |
| MoA                       | 0.18         | 0.12      | 0.26          | <u>0.34</u>  | <b>0.56</b>  |
| Nuc-in-Cyto               | 0.91         | 0.79      | 0.88          | <u>0.96</u>  | <b>0.97</b>  |
| Roundness                 | -0.24        | -0.32     | -0.34         | <u>-0.26</u> | <b>-0.19</b> |
| NucSize                   | <u>-0.88</u> | -1.02     | -2.21         | -1.04        | <b>-0.38</b> |
| CytoSize                  | -0.97        | -1.59     | -1.09         | <u>-0.65</u> | <b>-0.41</b> |
| NucCount                  | -2.61        | -1.05     | -0.83         | <u>-0.53</u> | <b>-0.28</b> |
| CytoCount                 | -3.22        | -1.39     | -1.03         | <u>-0.68</u> | <b>-0.33</b> |
| Overall                   | -5.20        | -3.19     | -2.44         | <u>0.46</u>  | <b>3.15</b>  |
| <i>Generation Quality</i> |              |           |               |              |              |
| FID                       | 41.94        | 35.70     | <b>20.36</b>  | 24.01        | <u>23.19</u> |
| KID                       | 0.028        | 0.029     | <u>0.015</u>  | <b>0.014</b> | 0.016        |



**Fig. 4: Qualitative comparisons.** *CellFluxRL* generates more biologically-grounded images, better capturing drug-induced morphological changes. In these examples, Etoposide-induced cell rounding, Demecolcine-driven microtubule destabilization, and AZ138-associated cell shrinkage are all more faithfully reproduced, and cell density more closely matches the ground truth for Cisplatin. Test-time scaling (+TTS) further refines these predictions toward the target images.

## 5.2 Reinforcement Learning Leads to Better Results

We compare *CellFluxRL* against baselines on a suite of biological, structural, and morphological metrics, with results summarized in Table 1. *CellFluxRL* outperforms its pretrained counterpart *CellFlux* across all reward metrics, confirming that RL post-training successfully steers the generative model toward greater physical correctness. The overall reward improves from  $-2.44$  to  $0.46$ ,



**Fig. 5:** Test-time scaling by best-of- $N$  further improves generation quality. The sample achieving the highest overall (combined) reward is selected from  $N$  rollouts, and each individual reward is plotted. RL (orange) consistently exhibits better scaling than the base model (blue) across all rewards.


with gains spanning all three reward categories: biological function, structural constraints, and morphological statistics.

*CellFluxRL* rewards also surpass those of *PhenDiff* and *IMPA*. The improvements are particularly pronounced on MoA, where *CellFluxRL*’s reward of 0.34 substantially exceeds both *PhenDiff* (0.18) and *IMPA* (0.12). To give interpretable context to these numbers, the corresponding MoA classification accuracies—the fraction of generated images whose predicted mode of action matches the ground truth label—are 0.56, 0.53, 0.61, and 0.66 for *PhenDiff*, *IMPA*, *CellFlux*, and *CellFluxRL* respectively, indicating that *CellFluxRL* produces images whose morphological profiles more faithfully distinguish different drug classes. Applying best-of- $N$  selection with  $N = 4$  to *CellFluxRL* (+TTS) yields further improvements across all metrics, pushing the overall reward from 0.46 to 3.15. We analyze test-time scaling behavior in further detail in §5.3.

In terms of generative quality, while not a primary optimization objective, *CellFluxRL* achieves a state-of-the-art KID score and a competitive FID score compared to baselines. The quantitative gains are corroborated qualitatively in Figure 4. The baselines generate images that fail to reflect the expected biological response to each perturbation. *CellFluxRL* consistently corrects these failures, and test-time scaling pushes generations further toward the ground truth.

### 5.3 Test-time Scaling Further Enhances the Results

Figure 5 shows how reward scores scale with the number of candidate samples  $N$  for both the base model (*CellFlux*) and the RL-post-trained model (*CellFluxRL*). We generate  $N$  candidates and select the one with the highest overall reward, the same weighted combination of individual rewards used during RL post-training, then report the individual reward metrics of the selected samples. Both models exhibit monotonic improvements as  $N$  increases across all metrics. RL post-training improves the scaling efficiency: the *CellFluxRL* curve sits consistently above the base model curve across all metrics, meaning that for any fixed compute budget (i.e., any given  $N$ ), the RL-post-trained model achieves higher reward. Because each sample drawn from the improved distribution is more likely



**Fig. 6:** Sensitivity analysis on KL weight  $\beta$ . Each subplot shows reward sensitivity to  $\beta$  after RL post-training.

**Table 2:** Single-reward optimization with reinforcement learning. Each column shows a model optimized for one reward only or *CellFluxRL* (all rewards), evaluated across the metrics in each row. **Bold** values indicate the best performance; underlined values indicate the second best; dashed-underlined values indicate the third best.

| Metric      | MoA          | Nuc-in-Cyto   | Roundness     | NucSize       | CytoSize      | NucCount      | CytoCount     | <i>CellFluxRL</i> |
|-------------|--------------|---------------|---------------|---------------|---------------|---------------|---------------|-------------------|
| MoA         | <b>0.412</b> | 0.204         | 0.263         | 0.301         | 0.212         | 0.279         | 0.277         | <u>0.337</u>      |
| Nuc-in-Cyto | 0.852        | <b>0.977</b>  | 0.907         | 0.904         | <u>0.976</u>  | 0.867         | 0.937         | <u>0.956</u>      |
| Roundness   | -0.342       | -0.332        | <b>-0.198</b> | -0.298        | -0.340        | -0.271        | -0.304        | <u>-0.264</u>     |
| NucSize     | -3.319       | <u>-1.724</u> | -2.328        | <b>-0.702</b> | -1.784        | -2.909        | -2.853        | <u>-1.038</u>     |
| CytoSize    | -1.198       | <u>-0.545</u> | -0.919        | -0.917        | <b>-0.497</b> | -1.113        | -0.801        | <u>-0.652</u>     |
| NucCount    | -0.760       | -0.904        | -0.685        | -0.818        | -0.881        | <b>-0.476</b> | <u>-0.511</u> | <u>-0.528</u>     |
| CytoCount   | -0.922       | -1.074        | -0.905        | -0.946        | -1.063        | <u>-0.629</u> | <b>-0.580</b> | <u>-0.677</u>     |

to be physically correct, selection is more efficient. The effect is especially pronounced at small  $N$ , where the quality of the base distribution matters most.

#### 5.4 Sensitivity Analysis on Kullback–Leibler Divergence Weight

We study the sensitivity of final model performance on the KL divergence weight,  $\beta$ , varying it from 1.0 to 1.3. The results, presented in Figure 6, show a trade-off: structural and morphological correctness metrics favor a smaller KL weight, whereas biological function rewards slightly improve with a larger weight. This suggests that achieving higher biological accuracy requires more substantial model adjustments than improving structural or morphological fidelity. In practice,  $\beta$  acts as a knob for this trade-off: we default to  $\beta = 1.0$  to favor the structural and morphological constraints while still improving biological accuracy, and recommend a larger  $\beta$  when MoA consistency is the priority.

#### 5.5 Ablation Optimizing for a Single Reward

To isolate the effect of individual rewards, we conduct ablations optimizing for one reward at a time. As shown in Table 2, this targeted optimization significantly improves performance on the chosen metric, but the improvements on other metrics are limited. In contrast, *CellFluxRL*, optimized with the combined weighted sum of all seven rewards, achieves either the second- or third-best score

on every individual metric. This demonstrates that the combined reward strategy enables balanced, multi-faceted improvement across all biological, structural, and morphological criteria simultaneously.

## 6 Conclusion

In this work, we tackle a central challenge for virtual cells: encouraging generated images to be physically correct and biologically plausible. We accomplish this by introducing a suite of biologically-informed rewards, which serve a unified role—as evaluation metrics, as training signals for reinforcement learning, and as a tool for test-time scaling. We use these rewards to optimize the state-of-the-art perturbation model *CellFlux*, yielding *CellFluxRL*. Across all rewards, our approach consistently improves physical plausibility and maintains overall image quality over the base model. Together, these results evolve virtual cell generation from pixel-level realism toward biological groundedness.

*Limitations and Future Work.* Our biological rewards are manually engineered based on domain expertise. A promising future direction is leveraging Large Language Models to automatically translate scientific literature into executable reward functions, significantly lowering the barrier to entry for novel domains. In this work, our core contribution is the modular multi-reward RL pipeline itself. As new generative architectures and more efficient RL algorithms emerge, they can be seamlessly integrated into this framework to further advance physically grounded generation. Finally, while we demonstrate our RL pipeline on virtual cell modeling, our approach provides a domain-agnostic framework for constrained generation. It can be adapted to other scientific fields governed by strict, non-differentiable rules, such as material science or medical imaging.

*Acknowledgement.* This work was supported in part by ONR Grant N00014-22-1-2110, NSF Grant 2205084, and the Stanford Institute for Human-Centered Artificial Intelligence (HAI). EBF is a Biohub, San Francisco, Investigator. S.Y. is a Chan Zuckerberg Biohub — San Francisco Investigator. E.L. was supported by the Stanford Institute for Human-Centered AI.

## References

1. Bereket, M., Karaletsos, T.: Modelling cellular perturbations with the sparse additive mechanism shift variational autoencoder. In: Oh, A., Naumann, T., Globerson, A., Saenko, K., Hardt, M., Levine, S. (eds.) *Advances in Neural Information Processing Systems*. vol. 36, pp. 1–12 (2023)
2. Black, K., Janner, M., Du, Y., Kostrikov, I., Levine, S.: Training diffusion models with reinforcement learning. *arXiv preprint arXiv:2305.13301* (2023)
3. Bourou, A., Boyer, T., Gheisari, M., Daupin, K., Dubreuil, V., De Thonel, A., Mezger, V., Genovesio, A.: Phendiff: Revealing subtle phenotypes with diffusion models in real images. In: *MICCAI* (2024)
4. Bunne, C., Roohani, Y., Rosen, Y., Gupta, A., Zhang, X., Roed, M., Alexandrov, T., AlQuraishi, M., Brennan, P., Burkhardt, D.B., et al.: How to build the virtual cell with artificial intelligence: Priorities and opportunities. *Cell* (2024)
5. Cai, Y., Li, K., Jia, M., Wang, J., Sun, J., Liang, F., Chen, W., Juefei-Xu, F., Wang, C., Thabet, A., Dai, X., Ju, X., Yuille, A., Hou, J.: Phygdpo: Physics-aware groupwise direct preference optimization for physically consistent text-to-video generation (2026), <https://arxiv.org/abs/2512.24551>
6. Caie, P.D., Walls, R.E., Ingleston-Orme, A., Daya, S., Houslay, T., Eagle, R., Roberts, M.E., Carragher, N.O.: High-content phenotypic profiling of drug response signatures across distinct cancer cells. *Molecular Cancer Therapeutics* (2010)
7. Chandrasekaran, S.N., Ackerman, J., Alix, E., Ando, D.M., Arevalo, J., Bennion, M., Boisseau, N., Borowa, A., Boyd, J.D., Brino, L., et al.: Jump cell painting dataset: morphological impact of 136,000 chemical and genetic perturbations. *BioRxiv* pp. 2023–03 (2023)
8. Chandrasekaran, S.N., Cimini, B.A., Goodale, A., Miller, L., Kost-Alimova, M., Jamali, N., Doench, J.G., Fritchman, B., Skepner, A., Melanson, M., et al.: Three million images and morphological profiles of cells treated with matched chemical and genetic perturbations. *Nature Methods* pp. 1–8 (2024)
9. Chen, C., Yoon, J., Wu, Y.F., Ahn, S.: Transdreamer: Reinforcement learning with transformer world models. In: *Deep RL Workshop NeurIPS 2021* (2021)
10. Cobbe, K., Kosaraju, V., Bavarian, M., Chen, M., Jun, H., Kaiser, L., Plappert, M., Tworek, J., Hilton, J., Nakano, R., et al.: Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168* (2021)
11. Fan, Y., Watkins, O., Du, Y., Liu, H., Ryu, M., Boutilier, C., Abbeel, P., Ghavamzadeh, M., Lee, K., Lee, K.: Dpok: Reinforcement learning for fine-tuning text-to-image diffusion models. *Advances in Neural Information Processing Systems* **36**, 79858–79885 (2023)
12. Fay, M.M., Kraus, O., Victors, M., Arumugam, L., Vuggumudi, K., Urbanik, J., Hansen, K., Celik, S., Cernek, N., Jagannathan, G., et al.: Rxrx3: Phenomics map of biology. *Biorxiv* pp. 2023–02 (2023)
13. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. In: *NeurIPS* (2020)
14. Huang, S., Wang, Z., Li, P., Jia, B., Liu, T., Zhu, Y., Liang, W., Zhu, S.C.: Diffusion-based generation, optimization, and planning in 3d scenes. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 16750–16761 (June 2023)
15. Hung, A.Z., Zhang, C.J., Sexton, J.Z., O’Meara, M.J., Welch, J.D.: Lumic: Latent diffusion for multiplexed images of cells. *bioRxiv* pp. 2024–11 (2024)

16. Johnson, G.T., Agmon, E., Akamatsu, M., Lundberg, E., Lyons, B., Ouyang, W., Quintero-Carmona, O.A., Riel-Mehan, M., Rafelski, S., Horwitz, R.: Building the next generation of virtual cells to understand cellular biology. *Biophysical Journal* (2023)
17. Kang, B., Yue, Y., Lu, R., Lin, Z., Zhao, Y., Wang, K., Huang, G., Feng, J.: How far is video generation from world model: A physical law perspective. In: *International Conference on Machine Learning*. pp. 28991–29017. PMLR (2025)
18. Lamiable, A., Champetier, T., Leonardi, F., Cohen, E., Sommer, P., Hardy, D., Argy, N., Massougboji, A., Del Nery, E., Cottrell, G., et al.: Revealing invisible cell phenotypes with conditional generative modeling. *Nature Communications* (2023)
19. Li, J., Cui, Y., Huang, T., Ma, Y., Fan, C., Yang, M., Zhong, Z.: Mixgrpo: Unlocking flow-based grpo efficiency with mixed ode-sde. *arXiv preprint arXiv:2507.21802* (2025)
20. Lightman, H., Kosaraju, V., Burda, Y., Edwards, H., Baker, B., Lee, T., Leike, J., Schulman, J., Sutskever, I., Cobbe, K.: Let’s verify step by step. In: *The twelfth international conference on learning representations* (2023)
21. Lin, H., Zala, A., Cho, J., Bansal, M.: Videodirectorgpt: Consistent multi-scene video generation via llm-guided planning. In: *COLM* (2024)
22. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: *The Eleventh International Conference on Learning Representations* (2023)
23. Lipman, Y., Havasi, M., Holderrieth, P., Shaul, N., Le, M., Karrer, B., Chen, R.T., Lopez-Paz, D., Ben-Hamu, H., Gat, I.: Flow matching guide and code. *arXiv preprint arXiv:2412.06264* (2024)
24. Liu, J., Liu, G., Liang, J., Li, Y., Liu, J., Wang, X., Wan, P., Zhang, D., Ouyang, W.: Flow-grpo: Training flow matching models via online rl. *arXiv preprint arXiv:2505.05470* (2025)
25. Liu, J., Liu, G., Liang, J., Yuan, Z., Liu, X., Zheng, M., Wu, X., Wang, Q., Qin, W., Xia, M., et al.: Improving video generation with human feedback. *arXiv preprint arXiv:2501.13918* (2025)
26. Liu, Q., Yin, X., Yuille, A., Brown, A., Singh, M.: Flowing from words to pixels: A framework for cross-modality evolution. *arXiv preprint arXiv:2412.15213* (2024)
27. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. In: *ICLR* (2023)
28. Ma, N., Tong, S., Jia, H., Hu, H., Su, Y.C., Zhang, M., Yang, X., Li, Y., Jaakkola, T., Jia, X., et al.: Inference-time scaling for diffusion models beyond scaling denoising steps. *arXiv preprint arXiv:2501.09732* (2025)
29. Matsuo, Y., LeCun, Y., Sahani, M., Precup, D., Silver, D., Sugiyama, M., Uchibe, E., Morimoto, J.: Deep learning, reinforcement learning, and world models. *Neural Networks* **152**, 267–275 (2022). <https://doi.org/https://doi.org/10.1016/j.neunet.2022.03.037>
30. Meng, F., Liao, J., Tan, X., Shao, W., Lu, Q., Zhang, K., Cheng, Y., Li, D., Qiao, Y., Luo, P.: Towards world simulator: Crafting physical commonsense-based benchmark for video generation. *arXiv preprint arXiv:2410.05363* (2024)
31. Pachitariu, M., Stringer, C.: Cellpose 2.0: how to train your own model. *Nature methods* **19**(12), 1634–1641 (2022)
32. Palma, A., Theis, F.J., Lotfollahi, M.: Predicting cell morphological responses to perturbations using generative modeling. *Nature Communications* (2025)
33. Qian, W., Wang, C., Peng, H., Tan, Z., Li, H., Zeng, A.: Rdp: Real data preference optimization for physics consistency video generation (2025), <https://arxiv.org/abs/2506.18655>



34. Rampášek, L., Hidru, D., Smirnov, P., Haibe-Kains, B., Goldenberg, A.: Dr.vae: improving drug response prediction via modeling of drug perturbation effects. *Bioinformatics* **35**(19), 3743–3751 (03 2019). <https://doi.org/10.1093/bioinformatics/btz158>
35. Slepchenko, B.M., Loew, L.M.: Use of virtual cell in studies of cellular dynamics. *International review of cell and molecular biology* **283**, 1–56 (2010)
36. Slepchenko, B.M., Schaff, J.C., Macara, I., Loew, L.M.: Quantitative cell biology with the virtual cell. *Trends in cell biology* (2003)
37. Snell, C., Lee, J., Xu, K., Kumar, A.: Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314* (2024)
38. Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., Ganguli, S.: Deep unsupervised learning using nonequilibrium thermodynamics. In: *ICML* (2015)
39. Song, Y., Ermon, S.: Generative modeling by estimating gradients of the data distribution. In: *NeurIPS* (2019)
40. Walker, D.C., Southgate, J.: The virtual cell—a candidate co-ordinator for ‘middle-out’ modelling of biological systems. *Briefings in Bioinformatics* **10**(4), 450–461 (03 2009). <https://doi.org/10.1093/bib/bbp010>
41. Wu, P., Escontrela, A., Hafner, D., Abbeel, P., Goldberg, K.: Daydreamer: World models for physical robot learning. In: Liu, K., Kulic, D., Ichnowski, J. (eds.) *Proceedings of The 6th Conference on Robot Learning. Proceedings of Machine Learning Research*, vol. 205, pp. 2226–2240. PMLR (14–18 Dec 2023)
42. Xu, J., Huang, Y., Cheng, J., Yang, Y., Xu, J., Wang, Y., Duan, W., Yang, S., Jin, Q., Li, S., Teng, J., Yang, Z., Zheng, W., Liu, X., Ding, M., Zhang, X., Gu, X., Huang, S., Huang, M., Tang, J., Dong, Y.: Visionreward: Fine-grained multi-dimensional human preference learning for image and video generation (2024), <https://arxiv.org/abs/2412.21059>
43. 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)
44. Yang, Y., Jia, B., Zhi, P., Huang, S.: Physcene: Physically interactable 3d scene synthesis for embodied ai. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 16262–16272 (June 2024)
45. You, L.: Toward computational systems biology. *Cell Biochemistry and Biophysics* **40**(2), 167–184 (2004). <https://doi.org/10.1385/CBB:40:2:167>
46. Zhang, Y., Su, Y., Wang, C., Li, T., Wefers, Z., Nirschl, J., Burgess, J., Ding, D., Lozano, A., Lundberg, E., et al.: Cellflux: Simulating cellular morphology changes via flow matching. *arXiv preprint arXiv:2502.09775* (2025)
47. Zheng, K., Chen, H., Ye, H., Wang, H., Zhang, Q., Jiang, K., Su, H., Ermon, S., Zhu, J., Liu, M.Y.: Diffusionnft: Online diffusion reinforcement with forward process. *arXiv preprint arXiv:2509.16117* (2025)
48. Zhu, H., He, T., Tang, A., Guo, J., Chen, Z., Bian, J.: Compositional 3d-aware video generation with llm director. In: Globerson, A., Mackey, L., Belgrave, D., Fan, A., Paquet, U., Tomczak, J., Zhang, C. (eds.) *Advances in Neural Information Processing Systems*. vol. 37, pp. 131618–131644 (2024). <https://doi.org/10.52202/079017-4184>