

# Reinforcement Learning for Generative AI: State of the Art, Opportunities and Open Research Challenges

**Giorgio Franceschelli**

GIORGIO.FRANCESCHELLI@UNIBO.IT

*Department of Computer Science and Engineering, University of Bologna, Bologna, Italy*

**Mirco Musolesi**

M.MUSOLESI@UCL.AC.UK

*Department of Computer Science, University College London, London, United Kingdom*

*Department of Computer Science and Engineering, University of Bologna, Bologna, Italy*

## Abstract

Generative Artificial Intelligence (AI) is one of the most exciting developments in Computer Science of the last decade. At the same time, Reinforcement Learning (RL) has emerged as a very successful paradigm for a variety of machine learning tasks. In this survey, we discuss the state of the art, opportunities and open research questions in applying RL to generative AI. In particular, we will discuss three types of applications, namely, RL as an alternative way for generation without specified objectives; as a way for generating outputs while concurrently maximizing an objective function; and, finally, as a way of embedding desired characteristics, which cannot be easily captured by means of an objective function, into the generative process. We conclude the survey with an in-depth discussion of the opportunities and challenges in this fascinating emerging area.

## 1. Introduction

Generative Artificial Intelligence (AI) is gaining increasing attention in academia, industry, and among the general public. This has been apparent since a portrait based on Generative Adversarial Networks (Goodfellow et al., 2014) was sold for more than four hundred thousand dollars (Christies, 2018) in 2018. Then, the introduction of transformers (Vaswani et al., 2017) for natural language processing and diffusion models (Sohl-Dickstein et al., 2015) for image generation has led to the development of generative models characterized by unprecedented performance, e.g., GPT-4 (OpenAI, 2023), LaMDA (Thoppilan et al., 2022), Llama 2 (Touvron et al., 2023), Gemini (Gemini Team & Google, 2023), DALL-E 2 (Ramesh et al., 2022) and Stable Diffusion (Rombach et al., 2022), just to name a few. In particular, ChatGPT (OpenAI, 2022), a conversational agent based on GPT-3 and GPT-4, is widely considered as a game-changing product; its introduction has indeed accelerated the development of foundation models. One of the characteristics of ChatGPT and other state-of-the-art large language models (LLMs) and foundation models<sup>1</sup> is the use of Reinforcement Learning (RL) in order to align its production to human values (Christiano et al., 2017), so as to mitigate biases and to avoid mistakes and potentially malicious uses.

---

1. We assume the following definitions: we refer to large language models as language models characterized by large size in terms of number of parameters; they are also usually based on transformer architectures. A foundation model is a large model that is trained on broad data of different types (textual, audio, image, video, etc.) at scale and is adaptable to a wide range of downstream tasks, following Bommasani et al. (2022).

In general, RL offers the opportunity to use non-differentiable functions as rewards (Ranzato et al., 2016). Examples include chemistry (Vanhaelen et al., 2020) and dialogue systems (Young et al., 2013). We believe that RL is a promising solution for designing efficient and effective generative AI systems. In this article, we will explore this research space, which is, after all, largely unexplored. In particular, the contributions of this work can be summarized as follows: we first survey the current state of the art at the interface (and intersection) between generative AI and RL. We systematize the existing literature according to three classes of applications, namely RL as an alternative way for generation with the goal of approximating outputs in the domain of interest as best as possible; as a way for generating outputs while concurrently maximizing quantifiable metrics or indicators; and, finally, as a way of embedding desired characteristics, which cannot be easily captured by means of an objective function, into the generative process. We then discuss the future opportunities and challenges of each category, outlining a potential research agenda for the coming years.

Several works have already surveyed deep generative learning (e.g., Franceschelli & Musolesi, 2021; Foster, 2023), deep reinforcement learning (e.g., Lazaridis et al., 2020; Sutton & Barto, 2018), its societal impacts (Whittlestone et al., 2021), and applications of RL for specific generative domains (e.g., Fernandes et al., 2023). To the best of our knowledge, this is the first survey on the applications (and implications) of RL applied to generative deep learning.

The remainder of the paper is structured as follows. First, we introduce and review key concepts in generative AI and RL (Section 2). Then, we discuss the different ways in which RL can be used for generative tasks, both considering past works and suggesting future directions (Section 3). Finally, we conclude the survey by discussing open research questions and analyzing future research opportunities (Section 4).

## 2. Preliminaries

### 2.1 Generative Deep Learning

We will assume the following definition of *generative model* (Foster, 2023): given a dataset of observations  $X$ , and assuming that  $X$  has been generated according to an unknown distribution  $P_{data}$ , a generative model  $P_{model}$  is a model that can mimic  $P_{data}$ . By sampling from  $P_{model}$ , observations that appear to have been drawn from  $P_{data}$  can be generated. Generative deep learning consists in the application of deep learning techniques to learn  $P_{model}$ .

Several families of generative deep learning techniques have been proposed in the last decade, e.g., Variational Autoencoders (VAEs; Kingma & Welling, 2014; Rezende et al., 2014), Generative Adversarial Networks (GANs; Goodfellow et al., 2014), autoregressive models like Recurrent Neural Networks (RNNs; Cho et al., 2014; Hochreiter & Schmidhuber, 1997), transformers (Vaswani et al., 2017), and denoising diffusion models (Sohl-Dickstein et al., 2015; Ho et al., 2020). These models and architectures aim to approximate  $P_{data}$  by means of self-supervised learning, i.e., by minimizing a reconstruction error when trying to reproduce real examples from  $X$ . The only exceptions are GANs, which aim to approximate  $P_{data}$  using adversarial learning, i.e., by maximizing the predicted probability that the outputs were generated by  $P_{data}$ . We refer the interested reader to Franceschelli

and Musolesi (2021) for a deeper analysis of the training and sampling processes at the basis of these solutions. Although highly effective for a variety of tasks, the outputs generated by these models do not always satisfy the desired properties. This happens for a variety of reasons. In fact, specific objectives cannot always be cast as loss functions; and providing carefully designed datasets is typically expensive. Few-shot learning (Brown et al., 2020), prompt engineering (Strobelt et al., 2023) and fine-tuning (Dodge et al., 2020) are potential solutions to these problems. We will discuss these issues in detail in the following sections.

## 2.2 Deep Reinforcement Learning




Figure 1: The canonical reinforcement learning framework: at each timestep  $t$ , the Agent performs an action  $a_t$  based on the current state  $s_t$ , which is a representation of the Environment. Upon the execution of the action, the Agent finds itself in a new state  $s_{t+1}$ , and receives a reward  $r_{t+1}$ .

RL is a machine learning paradigm that consists in learning an action based on a current representation of the environment in order to maximize a numerical signal, i.e., the *reward* over time (Sutton & Barto, 2018). More formally, at each time step  $t$ , an *agent* receives the current *state* from the *environment*, then it performs an *action* and observes the reward and the new state. Figure 1 summarizes the process. The learning process aims to teach the agent to act in order to maximize the *cumulative return*, i.e., a discounted sum of future rewards. Deep learning is also used to learn and approximate a *policy*, i.e., the mapping from states to action probabilities, or a *value function*, i.e., the mapping from states (or state-action pairs) to expected cumulative rewards. In this case, we refer to it as deep reinforcement learning. Several algorithms have been proposed to learn a value function from which it is possible to induce a policy, e.g., DQN (Mnih et al., 2013) and its variants (van Hasselt et al., 2016; Schaul et al., 2016; Wang et al., 2016), or to directly learn a policy, e.g., A3C (Mnih et al., 2016), DDPG (Lillicrap et al., 2016), TRPO (Schulman et al., 2015), PPO (Schulman et al., 2017). We refer the interested readers to Sutton and Barto (2018) for a comprehensive introduction to the topic.

The RL community has developed a variety of solutions to address the specific theoretical and practical problems emerging from this simple formulation. For example, if the reward signal is not known, inverse reinforcement learning (IRL; Ng & Russell, 2000) is used to

learn it from observed experience. Intrinsic motivation (Singh et al., 2004; Linke et al., 2020), e.g., curiosity (Pathak et al., 2017) can be used to deal with sparse rewards and encourage the agent to explore more. Imagination-based RL (Ha & Schmidhuber, 2018; Hafner et al., 2020) is a solution that allows to train an agent, reducing at the same time the need for interaction with the environment. Hierarchical RL (Pateria et al., 2021) allows to manage more complex problems by decomposing them into sub-tasks and working at different levels of abstraction. RL is not only used for training a single agent, but also in multi-agent scenarios (Zhang et al., 2021).

### 3. Reinforcement Learning for Generative AI

In the following, we will discuss the state of the art in RL for generative learning considering three classes of solutions, which are summarized in Table 1: RL as an alternative solution for output generation with the goal of approximating outputs from a given domain of interest with high fidelity; RL as a way for generating output while maximizing an objective function which captures (additional) quantifiable properties or indicators at the same time; and, finally, RL as a way of embedding additional desired characteristics (such as value alignment) which cannot easily be captured by means of an objective function into the generative process.

#### 3.1 Reinforcement Learning for Mere Generation

##### 3.1.1 OVERVIEW

The simplest approach is RL for *mere* generation, i.e., to train a generative model with the goal of approximating outputs from a given domain of interest as best as possible. Essentially, the objective function is used to replicate the behavior of the self-supervised learning loss used in traditional generative learning approaches, as the adversarial ones. In fact, due to its adherence to the formal framework of Markov decision processes (Sutton & Barto, 2018), RL can be used as a solution to the generative modeling problem in the case of sequential tasks (Bachman & Precup, 2015), e.g., text generation or stroke painting. The generative model plays the role of the agent. The current version of the generated output represents the state. For example, actions model how the state can be modified, e.g., which token<sup>2</sup> to be appended or which change applied to a picture. Finally, the reward is an indicator of the “quality” in terms of the generation of the output. Figure 2 summarizes the entire process.

It is possible to identify three fundamental design aspects: the implementation of the agent itself, e.g., diffusion model or transformer; the definition of the dynamics of the system, i.e., the transition between a state to another; the choice of the reward structure. The first two depend on the task to be solved, e.g., music generation with LSTM composing one note after the other or painting with CNN superimposing subsequent strokes. The third one is instead responsible of the actual learning. While the reward can be structured so as to represent the classic supervised target, it also provides the designers with the opportunity

---

2. We use the term “token” to refer to any discrete element an unstructured data point can be broken into, independently the data source is in the form of text (e.g., Radford et al., 2019), music (e.g., Huang & Wu, 2016), image (e.g., Ramesh et al., 2021) and so on.

Goal	Reward	Advantages	Limitations
Mere generation	<ul style="list-style-type: none"> <li>• GAN’s discriminative signal</li> <li>• Log-likelihood of real or predicted targets</li> <li>• Constraint satisfaction</li> </ul>	<ul style="list-style-type: none"> <li>• Models domains defined by non-differentiable objectives</li> <li>• Adapts GAN to sequential tasks</li> <li>• Can implement RL strategies, e.g., hierarchical RL</li> </ul>	<ul style="list-style-type: none"> <li>• Learning without supervision is hard</li> <li>• Pre-training can prevent an appropriate exploration</li> </ul>
Objective maximization	<ul style="list-style-type: none"> <li>• Test-time metrics</li> <li>• Countable desired or undesired characteristics</li> <li>• Distance-based measures</li> <li>• Quantifiable properties</li> <li>• Output of ML algorithms</li> </ul>	<ul style="list-style-type: none"> <li>• Satisfies quantifiable requirements</li> <li>• Optimizes a generator from a specific domain towards desirable sub-domains</li> <li>• Reduces the gap between training and evaluation</li> </ul>	<ul style="list-style-type: none"> <li>• Not every desirable property is quantifiable or easy to get</li> <li>• Goodhart’s law</li> </ul>
Improving not easily quantifiable characteristics	<ul style="list-style-type: none"> <li>• Output of a model trained to reproduce human or AI feedback about non-quantifiable properties (e.g., helpfulness, appropriateness, creativity)</li> </ul>	<ul style="list-style-type: none"> <li>• Satisfies non-quantifiable requirements (for example, the alignment problem)</li> <li>• Requires preferences between candidates instead of defining a mathematical measure of desired property</li> </ul>	<ul style="list-style-type: none"> <li>• Get user preferences is expensive</li> <li>• Users might misbehave, disagree, or be biased</li> <li>• Reward modeling is difficult</li> <li>• Prone to jailbreaks out of alignment</li> </ul>

Table 1: Summary of the three purposes for using RL with generative AI, considering the used rewards, their advantages, and their limitations.

of using a more diverse and complex set of reward functions, especially non-differentiable ones (which cannot be used in supervised learning due to the impossibility of computing their gradient for backpropagation).

The first example we consider is SeqGAN (Yu et al., 2017). Typically, GANs cannot be used for sequential tasks because the discriminative signal, i.e., whether the input looks real or not, is only available after the sequence is completed. SeqGAN circumvents this problem by using RL, which allows to learn from rewards received further in the future as well. Indeed, SeqGAN exploits the discriminative signal as the actual reward. The approach itself is based on a very simple policy approximation algorithm, namely REINFORCE (Williams, 1992). A similar approach is also used in MaskGAN (Fedus et al., 2018), where the generator learns with in-filling (i.e., by masking out a certain amount of words and then using the generator to predict them) through actor-critic learning (Sutton, 1984). Notably, hierarchical RL can also be used: for example, LeakGAN (Guo et al., 2018) relies on a generator




Figure 2: The reinforcement learning framework for generative modeling: at each timestep  $t$ , the Generative Model (i.e., the Agent) generates an action  $a_t$  based on the current description of the generated output (i.e., current state)  $s_t$ , which updates the current description of the generated output to  $s_{t+1}$ , and receives a reward  $r_{t+1}$  related to it.

composed of a manager, which receives *leaked* information from the discriminator, and a worker, which relies on a goal vector as a conditional input from the manager. Since SeqGAN might produce very sparse rewards, alternative strategies have been proposed. Shi et al. (2018) suggest to replace the discriminator with a reward model learned with IRL on state-action pairs, so that the reward is available at each timestep (together with an entropy regularization term). A more complex state composed of a context embedding can also be used (Li et al., 2019). Instead, Li et al. (2017) is based on a variation of SeqGAN: it uses Monte Carlo methods to get rewards at each timestep. In addition, the authors also suggest to alternate RL with a “teacher”, i.e., the classic supervised training. This helps deal with tasks like text generation where the action space (i.e., the set of possible words or sub-words) is too large to be consistently explored using RL alone. Another solution to this problem is NLPO (Ramamurthy et al., 2023), which is a parameterized-masked extension of PPO (Schulman et al., 2017) that restricts the action space via top- $p$  sampling, i.e., by only considering the smallest possible set of actions whose probabilities have a sum greater than  $p$  (Holtzman et al., 2020). TrufLL (Martin et al., 2022) uses top- $p$  sampling as well; however, it restricts the action space by means of a pre-trained task-agnostic model *before* applying policy gradient with PPO. Similarly, ColdGAN (Scialom et al., 2020) forces the sampling of a SeqGAN-like generator to be close to the distribution modes by selecting actions with top- $p$  sampling and low temperature (Holtzman et al., 2020) and training the generator via importance sampling (Precup et al., 2000). Finally, Lamprier et al. (2022) propose to substitute a top- $p$  sampling strategy with a cooperative one based on the use of Monte Carlo Tree Search structure, which is evaluated by the discriminator; again, the generator is trained via importance sampling.

Another reason to use RL is to take advantage of its inherent properties. For example, GOLD (Pang & He, 2021) is an algorithm that substitutes self-supervised learning with off-policy RL and importance sampling. It uses real demonstrations, which are stored in a replay buffer; the reward corresponds to either the sum or the product of the action probabilities over the sampled trajectories, i.e., of each single real token according to the model. While it can be considered close to a self-supervised approach, off-policy RL with importance sampling allows up-weighting actions with high (cumulative) return and actions preferred by the current policy, encouraging to focus on in-distribution examples.

RL is also an effective solution for learning in domains in which a differentiable objective is difficult or impossible to define. RL-Duet (Jiang et al., 2020) is an algorithm for online accompaniment generation. Learning how to produce musical notes according to a given context is a complex task: RL-Duet first learns a reward model that considers both inter-part (i.e., with counterpart) and intra-part (i.e., on its own) harmonization. Such model is composed by an ensemble of networks trained to predict different portions of music sheets (with or without human counter-part, and with or without machine context). Then, the generative system is trained to maximize this reward by means of an actor-critic architecture with generalized advantage estimator (GAE; Schulman et al., 2016). CodeRL (Le et al., 2022) performs code generation through a pre-trained model and RL. In particular, the model is fine-tuned with policy gradient in order to maximize the probability of passing unit tests: it receives a (sparse) reward quantifying if (and how) the generated code has passed the test for the assigned task. In addition, a critic learns a (dense) signal to predict the compiler output. The model is then trained to maximize both signals considering a baseline obtained with a greedy decoding strategy. In order to obtain a denser and more informative reward, PPOCoder (Shojaee et al., 2023) also considers three additional signals: a syntactic matching score based on the Abstract Syntax Tree of the generated code; a semantic matching score based on the data-flow graph; and a Kullback-Leibler (KL) penalty to prevent the model from deviating considerably from its pre-trained version. The sum of these four signals is then optimized via PPO.

Another interesting application area is painting. Xie et al. (2012) suggest to model stroke painting as a Markov Decision Process, where the state is the canvas, and the actions are the brushstrokes performed by the agent. Rewards calculated considering the location and inclination of the strokes are then used to train the agent. For instance, Doodle-SDQ (Zhou et al., 2018) fine-tunes a pre-trained sketcher with Double DQN (van Hasselt et al., 2016) and a reward that is calculated by evaluating how well a sketch reproduces a target image at pixel, movement, and color levels. Huang et al. (2019) use a discriminator trained to recognize real canvas-target image pairs to derive a corresponding reward. Instead, Singh and Zheng (2021) train a painting policy that operates at two different levels: foreground and background. Each of them uses a discriminator; in addition, they adopt a focus reward measuring the degree of indistinguishability of two object features. On the other hand, Intelli-Paint (Singh et al., 2022) is based on four different types of rewards, which are used to learn a painting policy with deep deterministic policy gradient (DDPG; Lillicrap et al., 2016) based on a discriminator signal on canvas-image pairs, two penalties for the color and position of consecutive strokes, and the same semantic guidance proposed by Singh and Zheng (2021). Finally, RL has also been used for collage artwork. Lee et al. (2023a) propose an RL-based method with the goal of composing different elements (such

as newspaper or texture cuts) in order to obtain an output that resembles a target picture. The state is composed by the canvas, the target image, and randomly (or value-based) sampled material; the action determines which region of the material to cut and where to paste it on current canvas; and the reward is the amount of similarity change between consecutive timesteps (where the similarity between the canvas and the target image is computed by a WGAN-GP discriminator (Gulrajani et al., 2017) trained in parallel to discriminate between target-target and target-canvas pairs). A model-based soft actor-critic (SAC; Haarnoja et al., 2019) is then used to optimize the reward minus a penalty for each timestep in order to teach the agent to complete the tasks with the minimum number of actions.

### 3.1.2 DISCUSSION

RL can represent an alternative method for deriving generative models, especially if the target loss is non-differentiable. It allows for the adaptation of known generative strategies, e.g., GANs, to tasks for which traditional techniques are not suitable, e.g., in text generation. In addition, it can be applied to domains in which feasibility and correctness (e.g., running code as above) are essential dimensions to consider. In other words, RL can train a generative model to produce observations that appear to have been drawn from the domain of interest even when such domain cannot be modeled by means of generative functions and corresponding differentiable losses. RL can also be used to derive more complex generative strategies (e.g., through hierarchical RL) and to reduce the model dependence on training data, which might have an impact on copyright issues (Franceschelli & Musolesi, 2022; Henderson et al., 2023).

It is possible to identify some limitations of the proposed solution. Learning without supervision is a hard task, especially when the reward is sparse. This is very likely to happen for sequence generation, such as (long) text or music, where the reward is available only at the last timestep. In addition to the aforementioned techniques for obtaining a denser reward, a potential solution might consist in considering an intrinsic reward (Aubret et al., 2019) as an additional learning signal, in order to encourage exploration as well. Moreover, the action space can be very large (potentially orders of magnitude larger than those of standard RL problems, Ammanabrolu & Hausknecht, 2020), especially for text generation. Ensuring a sufficient exploration of all possible actions while still exploiting the most promising ones to collect higher rewards is one of the key problems in RL. Starting with some prior knowledge about the possible best actions for different situations might be necessary for fast convergence. For this reason, pre-trained generative models are selected for this task. This can cause the agent to initially focus on highly probable tokens, increasing their associated probabilities and, because of that, failing to explore different solutions (i.e., by only moving the probability mass of the already most probable tokens) (Choshen et al., 2020). These problems can be avoided through variance reduction techniques (e.g., incorporating baselines and critics) and exploration strategies (Kiegeland & Kreutzer, 2021).



## 3.2 Reinforcement Learning for Objective Maximization

### 3.2.1 OVERVIEW

RL can be formalized and studied as an objective maximization problem. In this subsection we will discuss how this type of formalization can be applied to generative AI. Since RL allows us to use any non-differentiable function for modeling the rewards, it could be the case that simply replicating the behavior of a self-supervised learning loss is not the optimal solution. For example, Ranzato et al. (2016) point out the mismatch between how deep learning models are trained (i.e., on differentiable losses) and how they are commonly evaluated (i.e., on non-differentiable metrics): an emerging line of research is focusing on the use of non-differentiable metrics as reward functions for generative learning capturing a variety of requirements and constraints.

RL for quantity maximization has been mainly adopted in text generation, especially for dialogue and translation. In addition to exposure bias mitigation, it allows for replacing classic likelihood-based losses with metrics used at inference time. A pioneering work is the one by Ranzato et al. (2016), where RL is adopted to directly maximize BLEU (Papineni et al., 2002) and ROUGE (Lin, 2004) scores. To deal with the size of the action space, the authors introduce MIXER, a variant of REINFORCE algorithm that uses incremental learning (i.e., an algorithm based on an optimal pre-trained model according to ground truth sequences) and combines reward maximization with classic cross-entropy loss by means of an annealing schedule. In this way, the model starts with preexisting knowledge, which is preserved through the classic loss, while aiming at exploring alternative but still probable solutions, which should increase score at test time. A similar approach is also used by Google’s neural machine translation system (Wu et al., 2016). BLEU score is used as the reward, while fine-tuning a pre-trained neural translator with a mixed maximum likelihood and expected reward objective. Bahdanau et al. (2017) consider an actor-critic algorithm for machine translation, with the critic conditioned on the target text, and the pre-trained actor fine-tuned with BLEU as the reward. Paulus et al. (2018) suggest to learn to perform text summarization by using self-critical policy training (Rennie et al., 2017), where the reward associated with the action that would have been chosen at inference time is used as baseline. ROUGE score is considered as the reward, and linearly mixed with teacher forcing (Williams & Zipser, 1989), i.e., classic supervised learning. Scores alternative to ROUGE have been proposed as well, e.g., ROUGESal and Entail both described by Pasunuru and Bansal (2018). The former up-weights the salient sentences or words detected via a keyphrase classifier. The latter rewards logically-entailed summaries through an entailment classifier. They are then used alternatively in subsequent mini-batches to train a Seq2Seq model (Sutskever et al., 2014) by means of REINFORCE. Finally, Zhou et al. (2017) consider BLEU score to train a dialogue system on top of collected human interactions with offline RL. An additional dialogue-level reward function (measuring the number of proposed API calls) is also used. Recently, the RL4LM library (Ramamurthy et al., 2023) started offering many of these metrics as rewards, thus facilitating their use for LM training or fine-tuning. Different families of solutions are considered, i.e.,  $n$ -grams overlapping such as ROUGE, BLEU, SacreBLEU (Post, 2018) or METEOR (Lavie & Agarwal, 2007); model-based methods such as BertScore (Zhang et al., 2020) or BLEURT (Sellam et al., 2020);

task-specific metrics; and perplexity. Notably, RL4LM also allows to balance such metrics with a KL-divergence minimization with respect to a pre-trained model.

Test-time metrics are not the only quantities that can be maximized through RL. For example, Lagutin et al. (2021) suggest considering the count of 4-gram repetitions in the generated text, to reduce the likelihood of undesirable results. The combination of these techniques and classic self-supervised learning helps learn both *how to write* and *how not to write*. Li et al. (2016) train a Seq2Seq model for dialogue by rewarding conversations that are informative (i.e., which avoid repetitions), interactive (i.e., which reduce the probability of answers like “I don’t have any idea” that do not encourage further interactions), and coherent (i.e., which are characterized by high mutual information with respect to previous parts of the conversation). Sentence-level cohesion (i.e., compatibility of each pair of consecutive sentences) and paragraph-level coherence (i.e., compatibility among all sentences in a paragraph) can be achieved by maximizing the cosine similarity between the encoded version of the relative text, with the encoders trained so that the entire discriminative models are able to distinguish between real and generated pairs (Cho et al., 2019). A distance-based reward can instead guide a plot generator towards reaching desired goals. Tambwekar et al. (2019) train an agent working at event level (i.e., a tuple with the encoding of a verb, a subject, an object, and a fourth possible noun) with REINFORCE to minimize the distance between the generated verb and the goal verb. Other domain-specific rewards are used by Yi et al. (2018), where two distinct generative models produce poetry by maximizing fluency (i.e., MLE on a fixed language model), coherence (i.e., mutual information), meaningfulness (i.e., TF-IDF), and overall quality from a learned classifier. In addition, the two models also learn from each other: the worst performing can be trained on the output produced by the other one, or its distribution can be modified in order to better approximate the other.

Another popular technique is hierarchical RL: it allows optimization of quantifiable objectives even when they work at a different level of abstraction with respect to the generative model. For example, Peng et al. (2017) uses it to design a dialogue system able to perform composite tasks, i.e., sets of subtasks that need to be performed collectively. A high-level policy, trained to maximize an extrinsic reward directly provided by the user after each interaction, selects the sub-tasks. Then, “primitive” actions to complete the given sub-task are chosen according to a lower-level policy. A global state tracker on cross-subtask constraints is employed in order to provide the RL model with an intrinsic reward measuring how likely a particular subtask will be completed. Finally, ILQL (Snell et al., 2023) learns a state-action and a state-value function that are used to perturb a fixed LLM, rather than directly fine-tuning the model itself. This allows to preserve the capabilities of the given pre-trained language model, while still maximizing a specific utility function.

While text generation is one of the areas that have attracted most of the attention of researchers and practitioners in the past years, RL with quantity maximization has been applied to other sequential tasks as well. An important line of research (Jaques et al., 2016, 2017, 2017) consists of fine-tuning a pre-trained sequence predictor with imposed reward functions, while preserving the learned properties from data. For instance, a pre-trained note-based RNN can represent the starting point for the Q-network in DQN. A reward given by the probability of the chosen token according to the original model (or based on the inverse of the KL divergence) and one based on music theory rules (e.g., that all notes must belong to the same key) are used to fine-tune the model. Another possibility is to

extend SeqGAN to domain-specific reward maximization, as in ORGAN (Guimaraes et al., 2017). ORGAN linearly combines the discriminative signal with desired objectives, also dividing the reward by the number of repetitions made, in order to increase diversity in the result. Music generation can then be performed by considering tonality and ratio of steps as rewards; solubility, synthesizability and drug-likenesses are instead adopted to perform molecule generation as a sequential task, i.e., by considering a string-based representation of molecules (by means of SMILES language, Weininger, 1988). While the original work considered RNN-based models, transformer architectures can be used as well (Li et al., 2022).

Molecular generation is indeed one of the most explored task at the intersection between RL and generative AI. While MolGAN (De Cao & Kipf, 2018) adapts ORGAN to graph-based generative models (Li et al., 2018) to directly produce molecular structures, the majority of research focuses on simplified molecular-input line-entry system (SMILES) textual notation (Weininger, 1988), so as to leverage the recent advancements in text generation. ReLeaSe (Popova et al., 2018) fine-tunes a pre-trained generator to maximize physical, biological, or chemical properties (learned by a reward model). Olivecrona et al. (2017) propose to fine-tune a pre-trained generator with REINFORCE so as to maximize a linear combination of a prior likelihood (to avoid catastrophic forgetting) and a user-defined scoring function (e.g., to match a provided query structure or to have predicted biological activity). REINVENT (Blaschke et al., 2020) also avoids to generate molecules the model already produced through a memory that keeps track of the good scaffoldings generated so far. Atance et al. (2022) adopt REINVENT for the graph-based deep generative model GRAPHINVENT (Mercado et al., 2021) in order to directly obtain molecules that maximize desired properties, e.g., pharmacological activity. Instead, GENTRL (Zavoronkov et al., 2019) generates kinase inhibitors relying on a variational autoencoder to reduce molecules to continuous latent vectors. Then REINFORCE is used to teach the decoder how to maximize three properties learned through self-organizing maps: activity of compounds against kinases; closeness to neurons associated with DDR1 inhibitors within the whole kinase map; and novelty of chemical structures. The average reward for the produced batch is assumed as a baseline to reduce variance. Notably, RL is used here for single-step generation (i.e., by means of a contextual bandit). Gaudin et al. (2019) propose to generate molecules maximizing their partition coefficient without any pre-training by working with a simplified language (Krenn et al., 2020); Thiede et al. (2022) suggest to use intrinsic rewards to better explore its solution space. GCPN (You et al., 2018) trains a graph-CNN to optimize domain-specific rewards and an adversarial loss (from a GAN-like discriminator) through PPO. Other tasks have been investigated as well. Nguyen et al. (2022) merge GAN and actor-critic in order to obtain a generator capable of producing 3D material microstructures with desired properties. Han et al. (2020) use DDPG to train an agent to design buildings (in terms of shape and position) so as to maximize several signals related to the performance and aesthetics of the generated block, e.g., solar exposure, collision, and number of buildings.

Finally, the use of techniques based on objective maximization can also be effective for image generation. Denoising Diffusion Policy Optimization (DDPO; Black et al., 2023) can train or fine-tune a denoising diffusion model to maximize a given reward. It considers the iterative denoising procedure as a Markov Decision Process of fixed length. The state

contains the conditional context, the timestep, and the current image; each action represents a denoising step; and the reward is only available for the termination state, when the final, denoised image is obtained. DDPO has therefore been used to learn how to generate images that are more compressed or uncompressed, by minimizing or maximizing JPEG compression; more aesthetically pleasing, by maximizing LAION score (Schuhmann, 2022); or more prompt-aligned, by maximizing the similarity between the embeddings of prompt and generated image description. Improving the aesthetics of the image while preserving the text-image alignment has also been done at the prompt level (Hao et al., 2023). A language model that given human input provides an optimized prompt can be trained with PPO to maximize both an aesthetic score (from an aesthetic predictor) and a relevance score (as CLIP embedding similarity) of the image generated from the given prompt.

### 3.2.2 DISCUSSION

Reinforcement learning for objective maximization opens up several new possibilities: generators can be adapted for particular domains or for specific problems; they can be built for tasks difficult to model through differentiable functions; and pre-trained models can be fine-tuned according to given requirements and specifications. Essentially, RL is not used only for mere generation, since it also allows more specific, goal-oriented generative modeling: instead of training a generator to produce *correct, reasonable examples* for the domain of interest, the goal is to derive *the best possible examples* according to some specific target functions. Any desired and quantifiable property can now be set as reward function, thus in a sense “teaching” a model how to achieve it. While research has focused its attention on sequential tasks like text or music generation, other domains might be considered as well. As shown by Zhavoronkov et al. (2019), tasks not requiring multiple generative steps can be performed simply by reducing the RL problem to a contextual bandit one. In this way, RL can be considered as a technique for specific sub-domains, in a manner similar to neural style transfer (Gatys et al., 2016) or prompt engineering (Liu & Chilton, 2022).

We can identify possible drawbacks as well. Reinforcement learning has typically a very high computational cost (Ceron & Castro, 2021), due to the number of iterations required to converge. In addition, certain desired properties (e.g. harmlessness or appropriateness) can be difficult to quantify, or the related measures can be expensive to compute, especially at run-time. This can lead to excessive computational time for training. While offline RL might alleviate this problem, it would require a collection of evaluated examples, thus eliminating the advantage of not needing a dataset and increasing the risk of exposure bias. Finally, a fundamental issue arises from using test-time metrics as objective functions: how should we evaluate the model we derive? In fact, according to the empirical Goodhart’s Law (Goodhart, 1975), “when a measure becomes a target, it ceases to be a good measure”. New metrics are then required, and a gap between training objective and test score might be inevitable.

### 3.3 Reinforcement Learning for Improving Not Easily Quantifiable Characteristics

#### 3.3.1 OVERVIEW

While test-time metrics as objectives reduce the gap between training and evaluation, they not always correlate with human judgment (Chaganty et al., 2018). In these cases, using such metrics would not help obtain the desired generative model. Moreover, there might be certain qualities that do not have a correspondent metric because they are subjective, difficult to define, or, simply, not quantifiable. Typically users only have an implicit understanding of the task objective, and, therefore, a suitable reward function is almost impossible to design: this problem is commonly referred to as the *agent alignment problem* (Leike et al., 2018).

One of the most promising directions is reward modeling, i.e., learning the reward function from interaction with the user and then optimizing the agent through RL over such function. In particular, Reinforcement Learning from Human Feedback (RLHF; Christiano et al., 2017) allows to use human feedback to guide policy gradient methods. A reward model is trained to associate a reward to a trajectory thanks to human preferences (so that the reward associated with the preferred trajectory is higher than that associated with the others). In parallel, a policy is trained by means of this signal using a policy gradient method, while the trajectories collected at inference time are used to obtain new human feedback to improve the model. Ziegler et al. (2019) apply RLHF to text continuation, e.g., to write positive continuations of text summaries. A pre-trained language model is used to sample text continuations, which are then evaluated by humans; a reward model is trained over such preferences; and finally, the policy is fine-tuned using KL-PPO (Schulman et al., 2017) in order to maximize the signal provided by the reward model. A KL penalty is used to prevent the policy moving too far from its original version. Notably, these three steps can be performed once (offline case) or multiple times (online case).

Similarly, Stiennon et al. (2020) use RLHF to perform text summarization. The following three steps are repeated one or more times: human feedback collection, during which for each sampled Reddit post different summaries are generated, and then human evaluators are asked to rank them; reward model training on such preferences; policy training with PPO with the goal of maximizing the signal from the reward model (still using a KL penalty). Wu et al. (2021) propose to summarize entire books with RLHF by means of recursive task decomposition, i.e., by first learning to summarize small sections of a book, then summarizing those summaries into higher-level summaries, and so on. In this way, the size of the texts to be summarized is smaller. This is more efficient in terms of generative modeling and human evaluation, since the samples to be judged are shorter. InstructGPT (Ouyang et al., 2022) fine-tunes GPT-3 (Brown et al., 2020) with RLHF so that it can follow written instructions. With respect to Stiennon et al. (2020), demonstrations of desired behavior are first collected from humans and used to fine-tune GPT-3 before actually performing RLHF. Then, a prompt is sampled and multiple model outputs are generated, with a human labeler ranking them. Such rankings are finally used to train the reward model. The latter is then utilized (together with a KL penalty) to train the actual RL model with PPO. In particular, this procedure is adopted in ChatGPT and GPT-4 (OpenAI, 2023), which are fine-tuned in order to be aligned with human judgment.

Although all these methods consider human feedback regarding the “best” output for a given input (with “best” generally meaning appropriate, factual, respectful, or qualitative), more specific or different criteria are also used. Bai et al. (2022a) consider human preferences for helpfulness and harmlessness. Sparrow (Glaese et al., 2022) is trained to be helpful, correct, and harmless, with the three criteria judged separately so that three more efficient rule-conditional reward models are learned. In addition, the model is trained to query the web for evidence supporting provided facts; and again RLHF is used to obtain human feedback about the appropriateness of the collected evidence. Finally, Pardinias et al. (2023) use RLHF to fine-tune GPT-2 to learn how to write *haikus* maximizing the relevance to the provided topic, self-consistency, creativity, form, and avoiding toxic content through human feedback. In addition to text, RLHF has been used to better align text-to-image generation with human preferences. After collecting user feedback about text-image alignment, a reward model is learned to approximate such feedback, and its output is used to weight the classic loss function of denoising diffusion models (Lee et al., 2023c). On the contrary, DPOK (Fan et al., 2023) directly applies online reinforcement learning for fine-tuning text-to-image diffusion models, which are optimized using a learned reward model from human feedback (Xu et al., 2023) and a KL regularization with respect to the pre-trained model.

While very effective, RLHF is not the only existing approach. When human ratings are available in advance for each piece of text, a reward model can be trained offline and then used to fine-tune an LLM (Böhm et al., 2019). Such a reward model can also be combined with classic MLE to effectively train a language model (Kreutzer et al., 2018) or used to prepend reward tokens to generated text, forming a replay buffer suitable for online, off-policy algorithms to unlearn undesirable properties (Lu et al., 2022). Alternatively, A-LoL (Baheti et al., 2023) adopts offline policy gradient with a single-action step assumption (i.e., the entire sequence is a single action) to optimize for pre-trained, sequence-level reward models; in order to improve learning efficiency, it filters out data points with negative advantages, with the critic based on a frozen reference LLM. Since human ratings might be inaccurate, Nguyen et al. (2017) suggest to simulate them by applying perturbations on automatically generated scores. Alternatively, the provided dataset of scored text allows for batch (i.e., offline) policy gradient methods to train a chatbot (Kandasamy et al., 2017). A very similar approach is also followed by Jaques et al. (2020), where offline RL is used to train a dialogue system on collected conversations (with relative ratings) filtered to avoid learning misbehavior. Other strategies can be implemented as well. RELIS (Gao et al., 2019) relies on a learned reward model from human-provided judgment as the other systems discussed above; however, such reward model is used to optimize a policy directly at inference time for the provided text. Instead of training a policy over multiple inputs and then exploiting it at inference time, it trains a different policy for each required input.

Another possibility is to use AI feedback instead of, or in addition to, the human one. Constitutional AI (Bai et al., 2022b) is a method to train a non-evasive and “harmless” AI assistant without any human feedback, only relying on a *constitution* of principles to follow<sup>3</sup>.

---

3. While the selection of precepts to be included in the original “constitution” is defined by the researchers at Anthropic, a follow-up project called Collective Constitutional AI (Anthropic, 2023) involves the participation of humans for crowd-sourcing the underlying principles by means of Polis (Small et al., 2021), which is a platform for running online deliberative processes augmented by machine learning algorithms.

In a first supervised stage, a pre-trained LLM is used to generate responses to prompts, and then to iteratively correct them to satisfy a set of principles; once the response is deemed acceptable, it is used to fine-tune the model. Then, RLHF is performed, with the only difference that feedback is provided by the model itself and not by humans. RLAIIF (Lee et al., 2023b) completely replaces human preferences with preferences from an off-the-shelf LLM for text summarization. The desired overall behavior is induced by careful prompting. Liu et al. (2022) use RL to fine-tune a Seq2Seq model to generate knowledge useful for a generic QA model. This is first re-trained on knowledge generated with GPT-3 (which is prompted asking to provide the knowledge required to answer a certain question). Then, RL is used to fine-tune the model so as to maximize an accuracy score using knowledge generated by the model itself as a prompt. To avoid catastrophic forgetting, a KL penalty (with respect to the initial model) is introduced. RNEs (Wu & Hu, 2018) is instead a method to train an extractive summarizer (i.e., a component that selects which sentences of a given text should be included in its summary) using a reward based on coherence. A model is trained to identify the appropriate next sentence composing a coherent sentence pair; then, such a signal is used to obtain immediate rewards while training the agent (with ROUGE as the reward for the final composition). Finally, Su et al. (2016) propose to limit requests for human feedback to cases in which the learned reward model is uncertain.

### 3.3.2 DISCUSSION

Reward modeling introduces a great level of flexibility in RL for generative AI. Generative models can be trained to produce content that humans consider appropriate and of sufficient quality, by aligning them with their preferences. This is useful and in many situations essential: in fact a quantifiable measure might not exist or information to derive it might be hard to obtain. This methodology has already shown its intrinsic value in obtaining accurate, helpful, and useful text. In the same way, these techniques can be applied to other domains in which desired qualities are difficult to quantify or hard to express in a mathematical form, e.g., aesthetically pleasant or personalized (multimodal) content or creative artifacts (Franceschelli & Musolesi, 2023). A summary on the applications discussed in this paper is reported in Table 2.

RLHF has proven to be a highly effective approach. However, it suffers from several open problems (Casper et al., 2023). For example, getting user feedback can be incredibly expensive. Moreover, the users might misbehave, whether on purpose or not, be biased, or disagree within each other (Fernandes et al., 2023). Also, they might not correctly represent the population of end users or marginalized categories; and comparison-based feedback may not correlate with the desirability of responses (Casper et al., 2023). For these reasons, other techniques for modeling preferences might be considered. If human ratings are available in advance, a reward model can be derived from them and used in offline mode. Using AI itself to provide feedback is also an option; notably, AI-based feedback is also used outside the RL paradigm, e.g., to provide verbal feedback to be appended to prompts (Shinn et al., 2023) or to collaborate with other LLMs at inference time (Dong et al., 2023; Du et al., 2023). In addition, other techniques such as IRL or cooperative IRL (Hadfield-Menell et al., 2016) can be applied to induce a reward model from human demonstrations.

Application	Reward	RL Type	Example
<b>Chemistry</b>			
Molecule (graph)	Discriminator + chemical properties	P	(De Cao & Kipf, 2018)
	Pharmacological activity + prior likelihood	P	(Atance et al., 2022)
	Adversarial loss + desired properties	P	(You et al., 2018)
	Novelty + utility of inhibitors	CB	(Zavoronkov et al., 2019)
Molecule (text)	Discriminator + chemical properties	P	(Guimaraes et al., 2017)
	Learned desired properties	P	(Popova et al., 2018)
	Desired property + prior likelihood	P	(Olivecrona et al., 2017)
	As above + penalty for repetitions	P	(Blaschke et al., 2020)
	Partition coefficient	TD	(Gaudin et al., 2019)
	Desired property + intrinsic reward	P	(Thiede et al., 2022)
<b>Computer Vision</b>			
Collage	Discriminator on canvas-target pairs + length penalty	P	(Lee et al., 2023a)
Image	Compression or aesthetic or prompt alignment	P	(Black et al., 2023)
Stroke painting	Pixel, movement, color reproduction	TD	(Zhou et al., 2018)
	Discriminator on canvas-target pairs	P	(Huang et al., 2019)
	Background vs foreground + focus	P	(Singh & Zheng, 2021)
	Two above + adjacent color/position	P	(Singh et al., 2022)
Text-to-image	RLHF on text-image alignment	RWCE	(Lee et al., 2023c)
	Learned reward model from human feedback	P	(Fan et al., 2023)
<b>Design</b>			
Building	Performance and aesthetic metrics	P	(Han et al., 2020)
Microstructure	Adversarial loss + target properties	P	(Nguyen et al., 2022)
<b>Music</b>			
Accompaniment	Log-likelihood for pre-trained models	P	(Jiang et al., 2020)
Music	Discriminator signal	P	(Yu et al., 2017)
	Music theory rules + log-likelihood for original model	TD	(Jaques et al., 2016)
	Discriminator signal + tonality + ratio of steps	P	(Guimaraes et al., 2017)
<b>Natural language</b>			
Chatbot	Discriminator signal at each $t$ through MC methods	P	(Li et al., 2017)
	Discriminator signal at each $t$ through IRL	P	(Shi et al., 2018)
	Repetitive or useless answer penalty + mutual information	P	(Li et al., 2016)
	Reward from user + likelihood of sub-task completion	HP	(Peng et al., 2017)
	BLEU + number of proposed API calls	OffP	(Zhou et al., 2017)
	RLHF	P	(Ouyang et al., 2022)
	RLHF on helpfulness and harmlessness	P	(Bai et al., 2022a)
	RLHF on helpfulness, harmlessness and correctness	P	(Glaese et al., 2022)
	AI feedback based on a constitution of principles	P	(Bai et al., 2022b)
	Collected human ratings	OffP	(Kandasamy et al., 2017)
	Learned reward model of human ratings	TD	(Su et al., 2016)
	Learned sequence-level reward model of human preferences	OffP	(Baheti et al., 2023)
Extractive summarization	Reward model from human ratings	TD	(Gao et al., 2019)
	Coherence ratings + ROUGE	P	(Wu & Hu, 2018)
Generic text	Discriminator signal	P	(Fedus et al., 2018)
	Sum or product of log-likelihood of tokens from target text	OffP	(Pang & He, 2021)
	4gram repetition penalty + log-likelihood of target output	P	(Lagutin et al., 2021)
	Discriminator signals on coherence and cohesion	P	(Cho et al., 2019)
	Specific utility function to maximize at inference time	TD	(Snell et al., 2023)
Knowledge	Accuracy score + kl penalty	P	(Liu et al., 2022)
Machine translation	BLEU + log-likelihood of target output	P	(Ranzato et al., 2016)
	BLEU	P	(Bahdanau et al., 2017)
	Implicit task-based feedback from users	P	(Kreutzer et al., 2018)
	Perturbed predicted human ratings	CB	(Nguyen et al., 2017)
Plot	Generated vs target verbs distance	P	(Tambwekar et al., 2019)
Prompt optimization	Aesthetic score + CLIP similarity	P	(Hao et al., 2023)
Poetry	Discriminator signal	P	(Yu et al., 2017)
	Fluency + coherence + meaningfulness + quality	P	(Yi et al., 2018)
	RLHF on relevance, consistency, creativity, form, toxicity	P	(Pardinas et al., 2023)
	RLHF	P	(Ziegler et al., 2019)
Text continuation	ROUGE + log-likelihood of target output	P	(Paulus et al., 2018)
	ROUGESal + Entail	P	(Pasunuru & Bansal, 2018)
Text summarization	RLHF	P	(Stiennon et al., 2020)
	Reward model trained on human ratings	TD	(Böhm et al., 2019)
	RLAIF	P	(Lee et al., 2023b)
<b>Programming</b>			
Code	Result of unit tests	P	(Le et al., 2022)

Table 2: Summary of all the applications covered by past research in RL for generative AI, with the considered rewards and the relative references. Type of algorithms used: **On-Policy**; **Off-Policy**; **Temporal-Difference**; **Contextual Bandit**; **Hierarchical Policy**; **Reward-Weighted Cross-Entropy**.

Reward modeling can be problematic as well. Reducing the diversity of society to a single reward function might cause the majority views to disproportionately prevail (Feffer et al., 2023). In addition, seemingly well-performing preference-based reward models might



fail to generalize to out-of-distribution states (Tien et al., 2023), thus being prone to reward hacking (i.e., optimizing an imperfect proxy reward function that leads to poor performance according to the true reward function, Skalse et al., 2022). For these reasons, recent work has focused on eliminating the need for a reward model at all (e.g., Rafailov et al., 2023; Song et al., 2023).

Finally, Wolf et al. (2023) show that, even if aligned, LLMs can still be prompted in ways that lead to undesired behavior. In particular, “jailbreaks” out of alignment can be obtained via single prompts, especially when asking the model to simulate malicious personas (Deshpande et al., 2023). This is more likely to happen in the case of aligned models rather than of non-aligned ones because of the so-called *waluigi effect*: by learning to behave in a certain way, the model also learns its exact opposite (Nardo, 2023). More advanced approaches would be required to mitigate this problem and completely prevent certain undesired behaviors.

## 4. Conclusion

Reinforcement learning for generative AI has attracted huge attention after the recent breakthroughs in the area of foundation models and, in particular, large-scale language models. In this survey, we have investigated the state of the art, the opportunities and the open challenges in this fascinating area. First, we have discussed RL for classic generation, where RL simply provides a suitable framework for domains that cannot be modeled by means of a well-defined and differentiable objective, also reducing exposure bias. Then, we have considered RL for quantity maximization, where RL is used to teach a commonly pre-trained model how to maximize a numerical property. This closes the gap between what the model is optimized for and how it is evaluated, but also to search for particular characteristics and sub-domains. Finally, we have analyzed RL for non-easily quantifiable characteristics, where RL is used for aligning it with human requirements and preferences that are not easily expressed in a mathematical form.

Since non-differentiable functions can be used as target objectives, RL allows for a broader adoption of generative modeling, taking into consideration a wide range of objectives, requirements and constraints. Current and emerging solutions are characterized by the integration of a variety of RL mechanisms, such as IRL, hierarchical RL or intrinsic motivation, just to name a few. On the other hand, the use of RL for generative AI introduces the problem of balancing exploitation and exploration, especially when dealing with a large action space; this results in the need of using pre-trained models or a mixed objective both considering rewards and classic self-supervision. In addition, the adoption of test-time metrics as reward functions might be problematic per se (see the so-called Goodhart’s Law, Goodhart, 1975), while reward modeling is prone to human biases and adversarial attacks. Many challenging problems are still open, such as the integration of techniques such as IRL and multi-agent RL and the robustness of these models, in particular for preventing “jailbreaks” out of alignment.

## References

- Ammanabrolu, P., & Hausknecht, M. (2020). Graph constrained reinforcement learning for natural language action spaces. In *Proceedings of the 8th International Conference on Learning Representations (ICLR'20)*.
- Anthropic (2023). Collective Constitutional AI: Aligning a Language Model with Public Input. <https://anthropic.com/news/collective-constitutional-ai-aligning-a-language-model-with-public-input> [Accessed January 30, 2024].
- Atance, S. R., Diez, J. V., Engkvist, O., Olsson, S., & Mercado, R. (2022). De novo drug design using reinforcement learning with graph-based deep generative models. *Journal of Chemical Information and Modeling*, 62(20), 4863–4872.
- Aubret, A., Matignon, L., & Hassas, S. (2019). A survey on intrinsic motivation in reinforcement learning. arXiv:1908.06976 [cs.LG].
- Bachman, P., & Precup, D. (2015). Data generation as sequential decision making. In *Proceedings of the 28th International Conference on Neural Information Processing Systems (NIPS'15)*.
- Bahdanau, D., Brakel, P., Xu, K., Goyal, A., Lowe, R., Pineau, J., Courville, A., & Bengio, Y. (2017). An actor-critic algorithm for sequence prediction. In *Proceedings of the 5th International Conference on Learning Representations (ICLR'17)*.
- Baheti, A., Lu, X., Brahman, F., Bras, R. L., Sap, M., & Riedl, M. (2023). Improving Language Models with Advantage-based Offline Policy Gradients. arXiv:2305.14718 [cs.CL].
- Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., DasSarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., Joseph, N., Kadavath, S., Kernion, J., Conerly, T., El-Shawk, S., Elhage, N., Hatfield-Dodds, Z., Hernandez, D., Hume, T., ..., & Kaplan, J. (2022a). Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback. arXiv:2204.05862 [cs.CL].
- Bai, Y., Kadavath, S., Kundu, S., Askell, A., Kernion, J., Jones, A., Chen, A., Goldie, A., Mirhoseini, A., McKinnon, C., Chen, C., Olsson, C., Olah, C., Hernandez, D., Drain, D., Ganguli, D., Li, D., Tran-Johnson, E., Perez, E., ..., & Kaplan, J. (2022b). Constitutional AI: Harmlessness from AI Feedback. arXiv:2212.08073 [cs.CL].
- Black, K., Janner, M., Du, Y., Kostrikov, I., & Levine, S. (2023). Training diffusion models with reinforcement learning. In *ICML'23 Workshop on Efficient Systems for Foundation Models*.
- Blaschke, T., Arús-Pous, J., Chen, H., Margreitter, C., Tyrchan, C., Engkvist, O., Papadopoulos, K., & Patronov, A. (2020). REINVENT 2.0: An AI Tool for De Novo Drug Design. *Journal of Chemical Information and Modeling*, 60(12), 5918–5922.
- Böhm, F., Gao, Y., Meyer, C. M., Shapira, O., Dagan, I., & Gurevych, I. (2019). Better rewards yield better summaries: Learning to summarise without references. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP'19)*.

- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., ..., & Liang, P. (2022). On the Opportunities and Risks of Foundation Models. arXiv:2108.07258 [cs.LG].
- Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D., Wu, J., Winter, C., ..., & Amodei, D. (2020). Language models are few-shot learners. In *Advances in Neural Information Processing Systems (NIPS'20)*.
- Casper, S., Davies, X., Shi, C., Gilbert, T. K., Scheurer, J., Rando, J., Freedman, R., Korbak, T., Lindner, D., Freire, P., Wang, T., Marks, S., Segerie, C.-R., Carroll, M., Peng, A., Christoffersen, P., Damani, M., Slocum, S., Anwar, U., ..., & Hadfield-Menell, D. (2023). Open Problems and Fundamental Limitations of Reinforcement Learning from Human Feedback. arXiv:2307.15217 [cs.AI].
- Ceron, J. S. O., & Castro, P. S. (2021). Revisiting rainbow: Promoting more insightful and inclusive deep reinforcement learning research. In *Proceedings of the 38th International Conference on Machine Learning*.
- Chaganty, A., Mussmann, S., & Liang, P. (2018). The price of debiasing automatic metrics in natural language evaluation. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) (ACL'18)*, Melbourne, Australia.
- Cho, K., van Merriënboer, B., Bahdanau, D., & Bengio, Y. (2014). On the properties of neural machine translation: Encoder-decoder approaches. In *Proceedings of the 8th Workshop on Syntax, Semantics and Structure in Statistical Translation (SSST-8)*.
- Cho, W. S., Zhang, P., Zhang, Y., Li, X., Galley, M., Brockett, C., Wang, M., & Gao, J. (2019). Towards coherent and cohesive long-form text generation. In *Proceedings of the NAACL'19 Workshop on Narrative Understanding*.
- Choshen, L., Fox, L., Aizenbud, Z., & Abend, O. (2020). On the weaknesses of reinforcement learning for neural machine translation. In *Proceedings of the 8th International Conference on Learning Representations (ICLR'20)*.
- Christiano, P. F., Leike, J., Brown, T., Martic, M., Legg, S., & Amodei, D. (2017). Deep reinforcement learning from human preferences. In *Advances in Neural Information Processing Systems (NIPS'17)*.
- Christies (2018). Is artificial intelligence set to become art's next medium? <https://christies.com/stories/a-collaboration-between-two-artists-one-human-one-a-machine-0cd01f4e232f4279a525a446d60d4cd1> [Accessed January 12, 2024].
- De Cao, N., & Kipf, T. (2018). MolGAN: An implicit generative model for small molecular graphs. In *Proceedings of the ICML'18 Workshop on Theoretical Foundations and Applications of Deep Generative Models*.

- Deshpande, A., Murahari, V., Rajpurohit, T., Kalyan, A., & Narasimhan, K. (2023). Toxicity in chatgpt: Analyzing persona-assigned language models. In *Findings of the Association for Computational Linguistics: EMNLP 2023*.
- Dodge, J., Ilharco, G., Schwartz, R., Farhadi, A., Hajishirzi, H., & Smith, N. (2020). Fine-Tuning Pretrained Language Models: Weight Initializations, Data Orders, and Early Stopping. arXiv:2002.06305 [cs.CL].
- Dong, Y., Jiang, X., Jin, Z., & Li, G. (2023). Self-collaboration Code Generation via ChatGPT. arXiv:2304.07590 [cs.SE].
- Du, Y., Li, S., Torralba, A., Tenenbaum, J. B., & Mordatch, I. (2023). Improving Factuality and Reasoning in Language Models through Multiagent Debate. arXiv:2305.14325 [cs.CL].
- Fan, Y., Watkins, O., Du, Y., Liu, H., Ryu, M., Boutilier, C., Abbeel, P., Ghavamzadeh, M., Lee, K., & Lee, K. (2023). DPOK: Reinforcement Learning for Fine-tuning Text-to-Image Diffusion Models. arXiv:2305.16381 [cs.LG].
- Fedus, W., Goodfellow, I., & Dai, A. M. (2018). MaskGAN: Better Text Generation via Filling in the ----- . In *Proceedings of the 6th International Conference on Learning Representations (ICLR'18)*.
- Feffer, M., Heidari, H., & Lipton, Z. C. (2023). Moral machine or tyranny of the majority?. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence (AAAI'23)*.
- Fernandes, P., Madaan, A., Liu, E., Farinhas, A., Martins, P. H., Bertsch, A., de Souza, J. G. C., Zhou, S., Wu, T., Neubig, G., & Martins, A. F. T. (2023). Bridging the Gap: A Survey on Integrating (Human) Feedback for Natural Language Generation. arXiv:2305.00955 [cs.CL].
- Foster, D. (2023). *Generative Deep Learning*. O'Reilly.
- Franceschelli, G., & Musolesi, M. (2021). Creativity and Machine Learning: A Survey. arXiv:2104.02726 [cs.LG].
- Franceschelli, G., & Musolesi, M. (2022). Copyright in generative deep learning. *Data & Policy*, 4, e17.
- Franceschelli, G., & Musolesi, M. (2023). On the Creativity of Large Language Models. arXiv:2304.00008 [cs.AI].
- Gao, Y., Meyer, C. M., Mesgar, M., & Gurevych, I. (2019). Reward learning for efficient reinforcement learning in extractive document summarisation. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI'19)*.
- Gatys, L., Ecker, A., & Bethge, M. (2016). A neural algorithm of artistic style. *Journal of Vision*, 16(12), 326.
- Gaudin, T., Nigam, A., & Aspuru-Guzik, A. (2019). Exploring the chemical space without bias: data-free molecule generation with DQN and SELFIES. In *Second NIPS'19 Workshop on Machine Learning and the Physical Sciences*.
- Gemini Team, & Google (2023). Gemini: A Family of Highly Capable Multimodal Models. arXiv:2312.11805 [cs.CL].

- Glaese, A., McAleese, N., Trebacz, M., Aslanides, J., Firoiu, V., Ewalds, T., Rauh, M., Weidinger, L., Chadwick, M., Thacker, P., Campbell-Gillingham, L., Uesato, J., Huang, P.-S., Comanescu, R., Yang, F., See, A., Dathathri, S., Greig, R., Chen, C., ..., & Irving, G. (2022). Improving alignment of dialogue agents via targeted human judgments. arXiv:2209.14375 [cs.LG].
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., & Bengio, Y. (2014). Generative Adversarial Nets. In *Advances in Neural Information Processing Systems (NIPS'14)*.
- Goodhart, C. (1975). Problems of Monetary Management: The U.K. Experience. *Papers in Monetary Economics*, 1(1), 1–20.
- Guimaraes, G. L., Sanchez-Lengeling, B., Cunha Farias, P. L., & Aspuru-Guzik, A. (2017). Objective-Reinforced Generative Adversarial Networks (ORGAN) for Sequence Generation Models. arXiv:1705.10843 [stat.ML].
- Gulrajani, I., Ahmed, F., Arjovsky, M., Dumoulin, V., & Courville, A. C. (2017). Improved training of wasserstein gans. In *Advances in Neural Information Processing Systems (NIPS'17)*.
- Guo, J., Lu, S., Cai, H., Zhang, W., Yu, Y., & Wang, J. (2018). Long text generation via adversarial training with leaked information. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence and 30th Innovative Applications of Artificial Intelligence Conference and 8th AAAI Symposium on Educational Advances in Artificial Intelligence (AAAI'18/IAAI'18/EAAI'18)*.
- Ha, D., & Schmidhuber, J. (2018). Recurrent world models facilitate policy evolution. In *Advances in Neural Information Processing Systems (NIPS'18)*.
- Haarnoja, T., Zhou, A., Hartikainen, K., Tucker, G., Ha, S., Tan, J., Kumar, V., Zhu, H., Gupta, A., Abbeel, P., & Levine, S. (2019). Soft Actor-Critic Algorithms and Applications. arXiv:1812.05905 [cs.LG].
- Hadfield-Menell, D., Russell, S. J., Abbeel, P., & Dragan, A. (2016). Cooperative inverse reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS'16)*.
- Hafner, D., Lillicrap, T., Ba, J., & Norouzi, M. (2020). Dream to control: Learning behaviors by latent imagination. In *Proceedings of the 8th International Conference on Learning Representations (ICLR'20)*.
- Han, Z., Yan, W., & Liu, G. (2020). A performance-based urban block generative design using deep reinforcement learning and computer vision. In *Proceedings of 2020 DigitalFUTURES*.
- Hao, Y., Chi, Z., Dong, L., & Wei, F. (2023). Optimizing prompts for text-to-image generation. In *Proceedings of the 37th Conference on Neural Information Processing Systems (NIPS'23)*.
- Henderson, P., Li, X., Jurafsky, D., Hashimoto, T., Lemley, M. A., & Liang, P. (2023). Foundation Models and Fair Use. arXiv:2303.15715 [cs.CY].

- Ho, J., Jain, A., & Abbeel, P. (2020). Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems (NIPS'20)*.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural Computation*, 9(8), 1735–1780.
- Holtzman, A., Buys, J., Du, L., Forbes, M., & Choi, Y. (2020). The curious case of neural text degeneration. In *Proceedings of the 8th International Conference on Learning Representations (ICLR'20)*.
- Huang, A., & Wu, R. (2016). Deep Learning for Music. arXiv:1606.04930 [cs.LG].
- Huang, Z., Zhou, S., & Heng, W. (2019). Learning to paint with model-based deep reinforcement learning. In *Proceedings of the 2019 IEEE/CVF International Conference on Computer Vision (ICCV'19)*.
- Jaques, N., Gu, S., Bahdanau, D., Hernández-Lobato, J. M., Turner, R. E., & Eck, D. (2017). Sequence Tutor: Conservative Fine-Tuning of Sequence Generation Models with KL-control. In *Proceedings of the 34th International Conference on Machine Learning (ICML'17)*.
- Jaques, N., Gu, S., Turner, R. E., & Eck, D. (2016). Generating music by fine-tuning recurrent neural networks with reinforcement learning. In *NIPS'16 Workshop on Deep Reinforcement Learning*.
- Jaques, N., Gu, S., Turner, R. E., & Eck, D. (2017). Tuning recurrent neural networks with reinforcement learning. In *ICLR'17 Workshop*.
- Jaques, N., Shen, J. H., Ghandeharioun, A., Ferguson, C., Lapedriza, À., Jones, N., Gu, S., & Picard, R. W. (2020). Human-centric dialog training via offline reinforcement learning. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP'20)*.
- Jiang, N., Jin, S., Duan, Z., & Zhang, C. (2020). RL-Duet: Online Music Accompaniment Generation Using Deep Reinforcement Learning. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence, the 32nd Innovative Applications of Artificial Intelligence Conference, the 10th AAAI Symposium on Educational Advances in Artificial Intelligence (AAAI'20/IAAI'20/EAAI'20)*.
- Kandasamy, K., Bachrach, Y., Tomioka, R., Tarlow, D., & Carter, D. (2017). Batch policy gradient methods for improving neural conversation models. In *Proceedings of the 5th International Conference on Learning Representations (ICLR'17)*.
- Kiegeland, S., & Kreutzer, J. (2021). Revisiting the weaknesses of reinforcement learning for neural machine translation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL'21)*.
- Kingma, D. P., & Welling, M. (2014). Auto-encoding variational bayes. In *Proceedings of the 2nd International Conference on Learning Representations (ICLR'14)*.
- Krenn, M., Häse, F., Nigam, A., Friederich, P., & Aspuru-Guzik, A. (2020). Self-referencing embedded strings (SELFIES): A 100% robust molecular string representation. *Machine Learning: Science and Technology*, 1(4), 045024.

- Kreutzer, J., Khadivi, S., Matusov, E., & Riezler, S. (2018). Can neural machine translation be improved with user feedback?. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers) (NAACL'18)*.
- Lagutin, E., Gavrilov, D., & Kalaidin, P. (2021). Implicit unlikelihood training: Improving neural text generation with reinforcement learning. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume (EACL'21)*.
- Lamprier, S., Scialom, T., Chaffin, A., Claveau, V., Kijak, E., Staiano, J., & Piwowarski, B. (2022). Generative cooperative networks for natural language generation. In *Proceedings of the 39th International Conference on Machine Learning (ICML'22)*.
- Lavie, A., & Agarwal, A. (2007). Meteor: An automatic metric for mt evaluation with high levels of correlation with human judgments. In *Proceedings of the 2nd Workshop on Statistical Machine Translation (StatMT'07)*.
- Lazaridis, A., Fachantidis, A., & Vlahavas, I. (2020). Deep reinforcement learning: A state-of-the-art walkthrough. *Journal of Artificial Intelligence Research*, 69, 1421–1471.
- Le, H., Wang, Y., Gotmare, A. D., Savarese, S., & Hoi, S. C. H. (2022). CodeRL: Mastering Code Generation through Pretrained Models and Deep Reinforcement Learning. In *Advances in Neural Information Processing Systems (NIPS'22)*.
- Lee, G., Kim, M., Lee, Y., Lee, M., & Zhang, B.-T. (2023a). Neural collage transfer: Artistic reconstruction via material manipulation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV'23)*.
- Lee, H., Phatale, S., Mansoor, H., Lu, K., Mesnard, T., Bishop, C., Carbune, V., & Rastogi, A. (2023b). RLAIIF: Scaling Reinforcement Learning from Human Feedback with AI Feedback. arXiv:2309.00267 [cs.CL].
- Lee, K., Liu, H., Ryu, M., Watkins, O., Du, Y., Boutilier, C., Abbeel, P., Ghavamzadeh, M., & Gu, S. S. (2023c). Aligning Text-to-Image Models using Human Feedback. arXiv:2302.12192 [cs.LG].
- Leike, J., Krueger, D., Everitt, T., Martic, M., Maini, V., & Legg, S. (2018). Scalable agent alignment via reward modeling: a research direction. arXiv:1811.07871 [cs.LG].
- Li, C., Yamanaka, C., Kaitoh, K., & Yamanishi, Y. (2022). Transformer-based objective-reinforced generative adversarial network to generate desired molecules. In *Proceedings of the 31st International Joint Conference on Artificial Intelligence (IJCAI'22)*.
- Li, J., Monroe, W., Ritter, A., Jurafsky, D., Galley, M., & Gao, J. (2016). Deep reinforcement learning for dialogue generation. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing (EMNLP'16)*.
- Li, J., Monroe, W., Shi, T., Jean, S., Ritter, A., & Jurafsky, D. (2017). Adversarial learning for neural dialogue generation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP'17)*.
- Li, Y., Vinyals, O., Dyer, C., Pascanu, R., & Battaglia, P. (2018). Learning deep generative models of graphs. In *Proceedings of the 35th International Conference on Machine Learning (ICML'18)*.

- Li, Z., Kiseleva, J., & de Rijke, M. (2019). Dialogue generation: From imitation learning to inverse reinforcement learning. In *Proceedings of the 33rd AAAI Conference on Artificial Intelligence and 31st Innovative Applications of Artificial Intelligence Conference and 9th AAAI Symposium on Educational Advances in Artificial Intelligence (AAAI'19/IAAI'19/EAAI'19)*.
- Lillicrap, T. P., Hunt, J. J., Pritzel, A., Heess, N., Erez, T., Tassa, Y., Silver, D., & Wierstra, D. (2016). Continuous control with deep reinforcement learning. In *Proceedings of the 4th International Conference on Learning Representations (ICLR'16)*.
- Lin, C.-Y. (2004). ROUGE: A Package for Automatic Evaluation of Summaries. In *Proceedings of ACL'04 Workshop on Text Summarization Branches Out*.
- Linke, C., Ady, N. M., White, M., Degris, T., & White, A. (2020). Adapting behavior via intrinsic reward: A survey and empirical study. *Journal of Artificial Intelligence Research*, 69, 1287–1332.
- Liu, J., Hallinan, S., Lu, X., He, P., Welleck, S., Hajishirzi, H., & Choi, Y. (2022). Rainier: Reinforced knowledge introspector for commonsense question answering. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing (EMNLP'22)*.
- Liu, V., & Chilton, L. B. (2022). Design guidelines for prompt engineering text-to-image generative models. In *Proceedings of the 2022 CHI Conference on Human Factors in Computing Systems (CHI'22)*.
- Lu, X., Welleck, S., Hessel, J., Jiang, L., Qin, L., West, P., Ammanabrolu, P., & Choi, Y. (2022). QUARK: Controllable text generation with reinforced unlearning. In *Advances in Neural Information Processing Systems (NIPS'22)*.
- Martin, A., Quispe, G., Ollion, C., Le Corff, S., Strub, F., & Pietquin, O. (2022). Learning natural language generation with truncated reinforcement learning. In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL'22)*.
- Mercado, R., Rastemo, T., Lindelof, E., Klambauer, G., Engkvist, O., Chen, H., & Bjerrum, E. J. (2021). Graph networks for molecular design. *Machine Learning: Science and Technology*, 2(2), 025023.
- Mnih, V., Badia, A. P., Mirza, M., Graves, A., Harley, T., Lillicrap, T. P., Silver, D., & Kavukcuoglu, K. (2016). Asynchronous methods for deep reinforcement learning. In *Proceedings of the 33rd International Conference on Machine Learning (ICML'16)*.
- Mnih, V., Kavukcuoglu, K., Silver, D., Graves, A., Antonoglou, I., Wierstra, D., & Riedmiller, M. (2013). Playing atari with deep reinforcement learning. In *NIPS'16 Workshop on Deep Reinforcement Learning*.
- Nardo, C. (2023). The Waluigi Effect (mega-post). <https://lesswrong.com/posts/D7PumeYTDpfbTp3i7/the-waluigi-effect-mega-post> [Accessed January 30, 2024].
- Ng, A. Y., & Russell, S. J. (2000). Algorithms for inverse reinforcement learning. In *Proceedings of the 17th International Conference on Machine Learning (ICML'00)*.



- Nguyen, K., Daumé III, H., & Boyd-Graber, J. (2017). Reinforcement learning for bandit neural machine translation with simulated human feedback. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP'17)*, Copenhagen, Denmark.
- Nguyen, P. C. H., Vlassis, N. N., Bahmani, B., Sun, W., Udaykumar, H. S., & Baek, S. S. (2022). Synthesizing controlled microstructures of porous media using generative adversarial networks and reinforcement learning. *Scientific Reports*, 12(1), 9034–9049.
- Olivecrona, M., Blaschke, T., Engkvist, O., & Chen, H. (2017). Molecular de-novo design through deep reinforcement learning. *Journal of Cheminformatics*, 9(1), 48–61.
- OpenAI (2022). Introducing ChatGPT. <https://openai.com/blog/chatgpt> [Accessed January 12, 2024].
- OpenAI (2023). GPT-4 Technical Report. arXiv:2303.08774 [cs.CL].
- Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., Schulman, J., Hilton, J., Kelton, F., Miller, L., Simens, M., Askell, A., Welinder, P., Christiano, P. F., Leike, J., & Lowe, R. (2022). Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems (NIPS'22)*.
- Pang, R. Y., & He, H. (2021). Text generation by learning from demonstrations. In *Proceedings of the 9th International Conference on Learning Representations (ICLR'21)*.
- Papineni, K., Roukos, S., Ward, T., & Zhu, W.-J. (2002). BLEU: A Method for Automatic Evaluation of Machine Translation. In *Proceedings of the 40th Annual Meeting on Association for Computational Linguistics (ACL'02)*.
- Pardinas, R., Huang, G., Vazquez, D., & Piché, A. (2023). Leveraging human preferences to master poetry. In *Proceedings of the AAAI-23 Workshop on Creative AI Across Modalities*.
- Pasunuru, R., & Bansal, M. (2018). Multi-reward reinforced summarization with saliency and entailment. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers) (NAACL'18)*.
- Pateria, S., Subagdja, B., Tan, A.-h., & Quek, C. (2021). Hierarchical reinforcement learning: A comprehensive survey. *ACM Computing Surveys*, 54(5), 1–35.
- Pathak, D., Agrawal, P., Efros, A. A., & Darrell, T. (2017). Curiosity-driven exploration by self-supervised prediction. In *Proceedings of the 34th International Conference on Machine Learning (ICML'17)*.
- Paulus, R., Xiong, C., & Socher, R. (2018). A deep reinforced model for abstractive summarization. In *Proceedings of the 6th International Conference on Learning Representations (ICLR'18)*.
- Peng, B., Li, X., Li, L., Gao, J., Celikyilmaz, A., Lee, S., & Wong, K.-F. (2017). Composite task-completion dialogue policy learning via hierarchical deep reinforcement learning. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing (EMNLP'17)*.

- Popova, M., Isayev, O., & Tropsha, A. (2018). Deep reinforcement learning for de novo drug design. *Science Advances*, 4(7), eaap7885.
- Post, M. (2018). A Call for Clarity in Reporting BLEU Scores. In *Proceedings of the 3rd Conference on Machine Translation: Research Papers (WMT'18)*.
- Precup, D., Sutton, R. S., & Singh, S. P. (2000). Eligibility traces for off-policy policy evaluation. In *Proceedings of the 17th International Conference on Machine Learning (ICML'00)*.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., & Sutskever, I. (2019). Language Models are Unsupervised Multitask Learners. [https://cdn.openai.com/better-language-models/language\\_models\\_are\\_unsupervised\\_multitask\\_learners.pdf](https://cdn.openai.com/better-language-models/language_models_are_unsupervised_multitask_learners.pdf) [Accessed January 30, 2024].
- Rafailov, R., Sharma, A., Mitchell, E., Ermon, S., Manning, C. D., & Finn, C. (2023). Direct Preference Optimization: Your Language Model is Secretly a Reward Model. arXiv:2305.18290 [cs.LG].
- Ramamurthy, R., Ammanabrolu, P., Brantley, K., Hessel, J., Sifa, R., Bauckhage, C., Hajishirzi, H., & Choi, Y. (2023). Is reinforcement learning (not) for natural language processing: Benchmarks, baselines, and building blocks for natural language policy optimization. In *Proceedings of the 11th International Conference on Learning Representations (ICLR'23)*.
- Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., & Chen, M. (2022). Hierarchical Text-Conditional Image Generation with CLIP Latents. arXiv:2204.06125 [cs.CV].
- Ramesh, A., Pavlov, M., Goh, G., Gray, S., Voss, C., Radford, A., Chen, M., & Sutskever, I. (2021). Zero-shot text-to-image generation. In *Proceedings of the 38th International Conference on Machine Learning (ICML'21)*.
- Ranzato, M., Chopra, S., Auli, M., & Zaremba, W. (2016). Sequence level training with recurrent neural networks. In *Proceedings of the 4th International Conference on Learning Representations (ICLR'16)*.
- Rennie, S. J., Marcheret, E., Mroueh, Y., Ross, J., & Goel, V. (2017). Self-critical sequence training for image captioning. In *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR'17)*.
- Rezende, D. J., Mohamed, S., & Wierstra, D. (2014). Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the 31st International Conference on Machine Learning (ICML'14)*.
- Rombach, R., Blattmann, A., Lorenz, D., Esser, P., & Ommer, B. (2022). High-resolution image synthesis with latent diffusion models. In *Proceedings of the 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'22)*.
- Schaul, T., Quan, J., Antonoglou, I., & Silver, D. (2016). Prioritized experience replay. In *Proceedings of the 33rd International Conference on Machine Learning (ICML'16)*.
- Schuhmann, C. (2022). LAION-Aesthetics. <https://laion.ai/blog/laion-aesthetics/> [Accessed January 12, 2024].

- Schulman, J., Levine, S., Abbeel, P., Jordan, M., & Moritz, P. (2015). Trust region policy optimization. In *Proceedings of the 32nd International Conference on Machine Learning (ICML'15)*.
- Schulman, J., Moritz, P., Levine, S., Jordan, M. I., & Abbeel, P. (2016). High-dimensional continuous control using generalized advantage estimation. In *Proceedings of the 4th International Conference on Learning Representations (ICLR'16)*.
- Schulman, J., Wolski, F., Dhariwal, P., Radford, A., & Klimov, O. (2017). Proximal Policy Optimization Algorithms. arXiv:1707.06347 [cs.LG].
- Scialom, T., Dray, P.-A., Lamprier, S., Piwowarski, B., & Staiano, J. (2020). Coldgans: Taming language gans with cautious sampling strategies. In *Proceedings of the 34th International Conference on Neural Information Processing Systems (NIPS'20)*.
- Sellam, T., Das, D., & Parikh, A. (2020). BLEURT: Learning Robust Metrics for Text Generation. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics (ACL'20)*.
- Shi, Z., Chen, X., Qiu, X., & Huang, X. (2018). Toward diverse text generation with inverse reinforcement learning. In *Proceedings of the 27th International Joint Conference on Artificial Intelligence (IJCAI'18)*.
- Shinn, N., Cassano, F., Berman, E., Gopinath, A., Narasimhan, K., & Yao, S. (2023). Reflexion: Language Agents with Verbal Reinforcement Learning. arXiv:2303.11366 [cs.AI].
- Shojaee, P., Jain, A., Tipirneni, S., & Reddy, C. K. (2023). Execution-based Code Generation using Deep Reinforcement Learning. *Transactions on Machine Learning Research*.
- Singh, J., Smith, C., Echevarria, J., & Zheng, L. (2022). Intelli-Paint: Towards Developing More Human-Intelligible Painting Agents. In *Proceedings of the 17th European Conference on Computer Vision (ECCV'22)*.
- Singh, J., & Zheng, L. (2021). Combining semantic guidance and deep reinforcement learning for generating human level paintings. In *Proceedings of the 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR'21)*.
- Singh, S., Barto, A. G., & Chentanez, N. (2004). Intrinsically motivated reinforcement learning. In *Advances in Neural Information Processing Systems (NIPS'04)*.
- Skalse, J. M. V., Howe, N. H. R., Krasheninnikov, D., & Krueger, D. (2022). Defining and characterizing reward gaming. In *Advances in Neural Information Processing Systems (NIPS'22)*.
- Small, C., Bjorkegren, M., Erkkilä, T., Shaw, L., & Megill, C. (2021). Polis: Scaling deliberation by mapping high dimensional opinion spaces. *Recerca : Revista de Pensament i Anàlisi*, 26(2), 1–26.
- Snell, C. V., Kostrikov, I., Su, Y., Yang, S., & Levine, S. (2023). Offline RL for Natural Language Generation with Implicit Language Q Learning. In *Proceedings of the 11th International Conference on Learning Representations (ICLR'23)*.

- Sohl-Dickstein, J., Weiss, E., Maheswaranathan, N., & Ganguli, S. (2015). Deep unsupervised learning using nonequilibrium thermodynamics. In *Proceedings of the 32nd International Conference on Machine Learning (ICML'15)*.
- Song, F., Yu, B., Li, M., Yu, H., Huang, F., Li, Y., & Wang, H. (2023). Preference Ranking Optimization for Human Alignment. arXiv:2306.17492 [cs.CL].
- Stiennon, N., Ouyang, L., Wu, J., Ziegler, D., Lowe, R., Voss, C., Radford, A., Amodei, D., & Christiano, P. F. (2020). Learning to summarize with human feedback. In *Advances in Neural Information Processing Systems (NIPS'20)*.
- Strobelt, H., Webson, A., Sanh, V., Hoover, B., Beyer, J., Pfister, H., & Rush, A. M. (2023). Interactive and visual prompt engineering for ad-hoc task adaptation with large language models. *IEEE Transactions on Visualization and Computer Graphics*, 29(1), 1146–1156.
- Su, P.-H., Gašić, M., Mrkšić, N., Rojas-Barahona, L. M., Ultes, S., Vandyke, D., Wen, T.-H., & Young, S. (2016). On-line active reward learning for policy optimisation in spoken dialogue systems. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers) (ACL'16)*.
- Sutskever, I., Vinyals, O., & Le, Q. V. (2014). Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems (NIPS'14)*.
- Sutton, R. S. (1984). *Temporal Credit Assignment in Reinforcement Learning*. Ph.D. thesis, University of Massachusetts Amherst.
- Sutton, R. S., & Barto, A. G. (2018). *Reinforcement Learning: An Introduction*. The MIT Press.
- Tambwekar, P., Dhuliawala, M., Martin, L. J., Mehta, A., Harrison, B., & Riedl, M. O. (2019). Controllable neural story plot generation via reward shaping. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI'19)*.
- Thiede, L. A., Krenn, M., Nigam, A., & Aspuru-Guzik, A. (2022). Curiosity in exploring chemical spaces: intrinsic rewards for molecular reinforcement learning. *Machine Learning: Science and Technology*, 3(3), 035008.
- Thoppilan, R., De Freitas, D., Hall, J., Shazeer, N., Kulshreshtha, A., Cheng, H.-T., Jin, A., Bos, T., Baker, L., Du, Y., Li, Y., Lee, H., Zheng, H. S., Ghafouri, A., Menegali, M., Huang, Y., Krikun, M., Lepikhin, D., Qin, J., ..., & Le, Q. (2022). LaMDA: Language Models for Dialog Applications. arXiv:2201.08239 [cs.CL].
- Tien, J., He, J. Z.-Y., Erickson, Z., Dragan, A., & Brown, D. S. (2023). Causal confusion and reward misidentification in preference-based reward learning. In *Proceedings of the 11th International Conference on Learning Representations (ICLR'23)*.
- Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., Bikel, D., Blecher, L., Ferrer, C. C., Chen, M., Cucurull, G., Esiobu, D., Fernandes, J., Fu, J., Fu, W., ..., & Scialom, T. (2023). Llama 2: Open Foundation and Fine-Tuned Chat Models. arXiv:2307.09288 [cs.CL].
- van Hasselt, H., Guez, A., & Silver, D. (2016). Deep Reinforcement Learning with Double Q-Learning. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence (AAAI'16)*.

- Vanhaelen, Q., Lin, Y.-C., & Zhavoronkov, A. (2020). The advent of generative chemistry. *ACS Medicinal Chemistry Letters*, 11(8), 1496–1505.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, L., & Polosukhin, I. (2017). Attention is all you need. In *Advances in Neural Information Processing Systems (NIPS'17)*.
- Wang, Z., Schaul, T., Hessel, M., Van Hasselt, H., Lanctot, M., & De Freitas, N. (2016). Dueling network architectures for deep reinforcement learning. In *Proceedings of the 33rd International Conference on Machine Learning (ICML'16)*.
- Weininger, D. (1988). SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *Journal of Chemical Information and Computer Sciences*, 28(1), 31–36.
- Whittlestone, J., Arulkumaran, K., & Crosby, M. (2021). The societal implications of deep reinforcement learning. *Journal of Artificial Intelligence Research*, 70, 1003–1030.
- Williams, R. J. (1992). Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine Learning*, 8, 229–256.
- Williams, R. J., & Zipser, D. (1989). A learning algorithm for continually running fully recurrent neural networks. *Neural Computation*, 1(2), 270–280.
- Wolf, Y., Wies, N., Levine, Y., & Shashua, A. (2023). Fundamental Limitations of Alignment in Large Language Models. arXiv:2304.11082 [cs.CL].
- Wu, J., Ouyang, L., Ziegler, D. M., Stiennon, N., Lowe, R., Leike, J., & Christiano, P. (2021). Recursively Summarizing Books with Human Feedback. arXiv:2109.10862 [cs.CL].
- Wu, Y., Schuster, M., Chen, Z., Le, Q. V., Norouzi, M., Macherey, W., Krikun, M., Cao, Y., Gao, Q., Macherey, K., Klingner, J., Shah, A., Johnson, M., Liu, X., Kaiser, L., Gouws, S., Kato, Y., Kudo, T., Kazawa, H., ..., & Dean, J. (2016). Google’s Neural Machine Translation System: Bridging the Gap between Human and Machine Translation. arXiv:1609.08144 [cs.CL].
- Wu, Y., & Hu, B. (2018). Learning to extract coherent summary via deep reinforcement learning. In *Proceedings of the 32nd AAAI Conference on Artificial Intelligence and 30th Innovative Applications of Artificial Intelligence Conference and 8th AAAI Symposium on Educational Advances in Artificial Intelligence (AAAI'18/IAAI'18/EAAI'18)*.
- Xie, N., Hachiya, H., & Sugiyama, M. (2012). Artist agent: A reinforcement learning approach to automatic stroke generation in oriental ink painting. In *Proceedings of the 29th International Conference on Machine Learning (ICML'12)*.
- Xu, J., Liu, X., Wu, Y., Tong, Y., Li, Q., Ding, M., Tang, J., & Dong, Y. (2023). ImageReward: Learning and Evaluating Human Preferences for Text-to-Image Generation. arXiv:2304.05977 [cs.CV].
- Yi, X., Sun, M., Li, R., & Li, W. (2018). Automatic poetry generation with mutual reinforcement learning. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP'18)*.

- You, J., Liu, B., Ying, Z., Pande, V., & Leskovec, J. (2018). Graph convolutional policy network for goal-directed molecular graph generation. In *Advances in Neural Information Processing Systems (NIPS'18)*.
- Young, S., Gašić, M., Thomson, B., & Williams, J. D. (2013). POMDP-Based Statistical Spoken Dialog Systems: A Review. *Proceedings of the IEEE*, 101(5), 1160–1179.
- Yu, L., Zhang, W., Wang, J., & Yu, Y. (2017). SeqGAN: Sequence Generative Adversarial Nets with Policy Gradient. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI'17)*.
- Zhang, K., Yang, Z., & Başar, T. (2021). Multi-agent reinforcement learning: A selective overview of theories and algorithms. In *Handbook of Reinforcement Learning and Control*, pp. 321–384. Springer International Publishing.
- Zhang, T., Kishore, V., Wu, F., Weinberger, K. Q., & Artzi, Y. (2020). BERTScore: Evaluating Text Generation with BERT. In *Proceedings of the 8th International Conference on Learning Representations (ICLR'20)*.
- Zhavoronkov, A., Ivanenkov, Y. A., Aliper, A., Veselov, M. S., Aladinskiy, V. A., Aladin-skaya, A. V., Terentiev, V. A., Polykovskiy, D. A., Kuznetsov, M. D., Asadulaev, A., Volkov, Y., Zholus, A., Shayakhmetov, R. R., Zhebrak, A., Minaeva, L. I., Zagribelnyy, B. A., Lee, L. H., Soll, R., Madge, D., ..., & Aspuru-Guzik, A. (2019). Deep learning enables rapid identification of potent DDR1 kinase inhibitors. *Nature Biotechnology*, 37(9), 1038–1040.
- Zhou, L., Small, K., Rokhlenko, O., & Elkan, C. (2017). End-to-end offline goal-oriented dialog policy learning via policy gradient. In *Proceedings of the NIPS'17 Workshop on Conversational AI*.
- Zhou, T., Fang, C., Wang, Z., Yang, J., Kim, B., Chen, Z., Brandt, J., & Terzopoulos, D. (2018). Learning to Sketch with Deep Q Networks and Demonstrated Strokes. arXiv:1810.05977 [cs.CV].
- Ziegler, D. M., Stiennon, N., Wu, J., Brown, T. B., Radford, A., Amodei, D., Christiano, P., & Irving, G. (2019). Fine-Tuning Language Models from Human Preferences. arXiv:1909.08593 [cs.CL].