

Towards a Framework for Human-AI Interaction Patterns in Co-Creative GAN Applications

Imke Grabe, Miguel González-Duque, Sebastian Risi and Jichen Zhu

IT University of Copenhagen (ITU), Rued Langgaards Vej 7, 2300 Copenhagen, Denmark

Abstract

With the rise of Generative Adversarial Networks (GANs), AI has increasingly become a partner to human designers in co-creating cultural artifacts. While generative models have been applied in various creative tasks across disciplines, a theoretical foundation for understanding human-GAN collaboration is yet to be developed. Drawing from the mixed-initiative co-creation community, we propose a preliminary framework to analyze co-creative GAN applications. We identify four primary interaction patterns: *Curating*, *Exploring*, *Evolving*, and *Conditioning*. The suggested framework enables us to discuss the affordances and limitations of the different kind of interactions underlying co-creative GAN applications.

Keywords

generative adversarial networks, mixed-initiative co-creation, human-AI collaboration

1. Introduction

With the recent development in human-AI collaborative applications, deep generative models have become potential co-designers to humans in creative tasks. Especially *Generative Adversarial Networks* (GANs) have gained attention from designers and non-programmers across creative disciplines such as architecture, fashion, computer games, and art [1]. Over the recent years, GANs have been of particular interest to the human-computer interaction (HCI) community for their ability to create high-resolution output [2, 3, 4].

While technological advancements underlying GANs progress rapidly [5, 6, 7], how to develop GANs that can effectively co-create with human designers is still an open problem. Recent HCI research has identified various challenges for designers to work with machine learning (ML) models [8, 9]. More specifically, Buschek et al. [10] outline the challenges in designing interactions with generative models. A key difficulty when designing with GANs is the knowledge gap regarding the technical possibilities and limitations of the models [1]. Due to its complexity, designers who apply GANs in creative processes cannot oversee their technical functioning. At the same time, ML engineers who develop the algorithmic aspects of the models are not acquainted with the design requirements of a particular creative task.

Research in mixed-initiative co-creation (MI-CC) [11, 12] offers a promising starting point to bridge this knowledge gap. The research area of MI-CC focuses on how humans and machines can co-create together, including the

co-creation with AI agents [8]. For example, Spoto and Oleynik [13] propose a framework to map the primary interactions taken in the co-creative process between humans and computational agents [13]. Muller et al. [14] further expand the above-mentioned framework by including actions specific to generative AI interfaces. While Muller et al.'s expanded framework provides a tool for analyzing generative AI interfaces, its extensive action set is not aligned with how the functioning of latent variable models, like GANs' generators, integrates into co-creation. We argue that a smaller but more specific set of technically-grounded actions is sufficient to describe GANs' interactive capabilities accurately.


We hence suggest tailoring the framework to co-creative GAN applications by 1) reducing the action space to a minimal set and 2) incorporating new actions better aligned with the algorithmic properties of GANs. With this vocabulary, we aim to describe co-creative GAN applications while providing insight into GANs' inner workings. We developed and tested our framework by analyzing related literature in art, computer games, fashion, and object design.

This paper aims to answer the question: *How do co-creative GAN applications support co-creativity?* To do so, we adapt an existing MI-CC framework previously applied to analyze a broad range of co-creative interfaces [13], and further developed to describe generative models [14]. Based on the MI-CC frameworks, we suggest a preliminary framework for in-depth analysis of co-creative GAN applications. We show that the suggested taxonomy lets us analyze emerging patterns in the interaction with GANs. We limit our study to GANs because the majority of examples in the field of generative design use them, but we believe that our framework applies to other deep generative models. More specifically, the proposed set of actions supports the mapping

Joint Proceedings of the ACM IUI Workshops 2022, March 2022, Helsinki, Finland

✉ imgr@itu.dk (I. Grabe); migd@itu.dk (M. González-Duque); sebr@itu.dk (S. Risi); jicz@itu.dk (J. Zhu)

© 2022 Copyright © 2022 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

of interactions with generators via latent codes. Hence, these actions are applicable to other generative models, such as *variational autoencoders* (VAEs), consisting of a generator network with outputs that are sampled based on alterable latent codes.

By synthesizing existing GAN applications and MI-CC theory, our work aims to shed light on how GANs participate in co-creative design processes. With the framework's GAN-specific grounding, we hope to make the technical and thereby interactive capabilities better understandable when mapping interactions. By increasing awareness of the system's functioning, this insight might allow us, and non-experts in particular, to design better interactive applications. With the identification of emerging patterns, we aim to enable a discussion on the current trends in co-creative GAN applications and illuminate possible future research in this field.

In summary, this paper presents our preliminary framework for mapping the interaction in co-creative GAN applications and discusses the co-creative affordance of the emerging interaction patterns. The rest of the paper is structured as follows. In the next section, we summarize related work. Then, we present the framework supported by exemplary co-creative GAN applications drawing on existing literature. We discuss how the identified interaction patterns support co-creativity, followed by a conclusion and outlook on future work.

2. Related Work

In this section, we summarize theoretical work on understanding co-creative GAN applications and approaches to map MI-CC interactions.

2.1. Understanding human-AI collaboration with GANs

At the intersection of HCI and AI research, designing human-AI interaction remains an open issue due to the complexity of AI models and their lack of interpretability [9, 15]. This includes designing co-creative GAN applications. GANs learn to generate content from latent features, but this mapping is not explainable and black-box, making it difficult to control the desired output after training [5, 7]. As a generative model, GANs are a novel design tool that brings new challenges when it comes to interaction [10]. Though non-programmers and designers have applied the tool, a theoretical framework for analyzing GANs' co-creative potential is yet to be developed. While surveys have been conducted on the technical aspects [16, 17, 18, 19] as well as its application in creative domains [20], there has not been much work on understanding GANs' role in co-creative applications. Research in the area mainly consists of studies

proposing new GAN applications [21, 22, 23, 24] or suggesting algorithmic methods to control GANs without pre-defined use cases [25, 26, 27]. As part of a study on AI-augmented typeface design, Zeng et al. [28] frame a GAN as a source of inspiration in a creativity model. To date, the only (systematic) review in the area is Hughes et al.'s [1] survey on collaborative applications of GANs in design tasks. Their analysis focuses on the design domains, inspecting applications of GANs by discipline. Across all approaches, they identify beautification and variation as modes of operation. While this helps us understand *what* the GAN adds to the creative process, it does not reflect on *how* it is involved in co-creation. We aim to dig deeper into the mechanisms at play and the creative role of GANs in collaborative tasks.

2.2. Mapping mixed-initiative co-creation

The research area MI-CC investigates interactive processes in which both human and machine contribute proactively to producing artifacts [12]. MI-CC applications can be placed on a continuum, describing the extent to which human and computational agent take initiative in the creative process [29, 30]. To analyze the initiatives taken in such collaboration, several theoretical approaches have been suggested. Spoto and Oleynik [13] proposed to decompose the MI-CC process into seven primary actions: *ideate, constrain, produce, suggest, select, assess, adapt*. The actions can be used to map the "creative flow" [13] between human and computer as a graph. By mapping the series of actions taken by the two agents, the authors outline the creative (iterative) process for more than 70 works ranging from game level creation to manufacturing design systems. When it comes to including AI agents in co-creation with human designers, new interaction possibilities emerge, requiring comprehensive analysis frameworks. Muller et al. [14] adapt the notation of Spoto and Oleynik's framework to processes including generative AI agents. They add four actions that are specific to generative models, such as *learn* to describe the computer's training process. Their extended framework can be used to visualize "human-AI interaction patterns in the generative space" [14, p.1]. Our work aims to tailor these MI-CC methods to analyze co-creative GAN applications. We build on existing mapping methods suggested by MI-CC to outline common interaction patterns in co-creation between human and GAN.

3. Mapping interaction patterns with GANs (GAN-MIP)

We suggest a framework that applies mixed-initiative methods to understand how GANs function in co-

creation. As a starting point of this study, we reviewed GAN applications in salient design domains. Reviewed studies ranged from prompting GANs to invent new art styles [31], over breeding game levels [21], to adjusting fashion items as part of a person’s outfit [25]. To make sense of the surveyed GAN studies, we leveraged existing MI-CC frameworks to map the actions in co-creative tasks. They allowed us to analyze how GANs get involved in creative processes. In an iterative manner, we analyzed studies and revised the frameworks to derive one that is applicable across co-creative GAN applications. During this process, we focused on how the frameworks’ actions reflect the GANs’ technical functioning. Along the way, predominant flows in how humans and GANs interact in co-creation stood out among the studies, which we refer to as *interaction patterns*. This section presents the suggested framework and demonstrates the four primary interaction patterns we identified with it.

3.1. Framework

Our framework builds on the framework suggested by Spoto and Oleynik [13], further extended to map co-creation with generative models by Muller et al. [14].

We adapt their proposed actions to a minimal set required to describe co-creative GAN applications. For example, Muller et al. distinguish between *producing* one artifact and *suggesting* a set of artifacts to choose from as separate actions. We combine them into one category called *create*, as the artifact sampling from the GAN depicts the same action in either case. In addition, we add new actions, such as *initialize*, which we find can present a crucial design choice taken by human agents [24]. The final set of actions illustrates the co-creation as a graph mapping along the two axes of agents and actions. This section defines the agents making up the horizontal axis and explains what the actions on the vertical axis stand for.

3.1.1. Agents

The agents conducting the actions are (1) a human designer and (2) a computational system including a GAN as the main technology, which might be supported by other algorithms facilitating the interaction with it.

3.1.2. Actions

The actions are to be understood as activities that influence the artifact or the other agent directly. How actions can be executed depends on several factors, such as the design of the interface, the implementation details of a model, or the abilities of a human agent. Seven actions are used to map the interaction between the two agents, as described below.

Initialize. To initialize refers to setting up the GAN, such as choosing a dataset and a model architecture.

Learn. To learn describes the process of internalizing information. In line with Muller et al. [14], this refers to the training GANs on a chosen data distribution. For the human counterpart, the action would correspond to learning new skills or adapting to a new domain.

Constrain. To constrain refers to the process of specifying desired characteristics of the target artifacts, hence restricting the conceptual space.

Create. To create describes the action of generating new (candidates of) artifacts. The number of artifacts to be created can vary from one to many.

Select. To select refers to the action of choosing one or more artifacts and excluding others from the further process. The action encompasses different selection mechanisms including the structuring of subsets such as ranking.

Adapt. To adapt describes the act of making changes to existing artifacts. The agents can edit artifacts directly or adjust their representation, namely the latent code, that is used to *create* a new artifact.

Combine. To combine describes the process of constructing a new artifact that inherits parts of existing artifacts. As with adapting, this is done via the artifacts’ latent code.

The actions described above present a minimal set sufficient to describe the core interactions of co-creative GAN applications. For example, we left out the action *ideate*, which describes creating high-level concepts in existing frameworks [13, 14]. As we find that conceptual ideation is often expressed by conditions imposed on a GAN, the action *constrain* covers it in our framework, aligned with the GANs’ algorithmic function. The derived action set aims to align with the algorithmic properties of GANs as well as the interactive abilities of the human designer. Mapping interactions with the framework may help non-technical users better understand a co-creative partner’s algorithmic actions and show how the human designer’s actions embed into it. The set of core actions allows us to map the interactions of the co-creative GAN applications we surveyed.

3.2. Primary Interaction Patterns

Using the above framework of agents and actions, we identified four interaction patterns which we present along with an exemplary selection of co-creative GAN applications in which they appear. We reviewed studies that propose models for co-creation with GANs tested in user studies, such as a GAN that allows game designers to generate Mario levels [21]. Besides that, we also considered approaches that are not explicitly embedded into use cases yet but suggest applicable models that

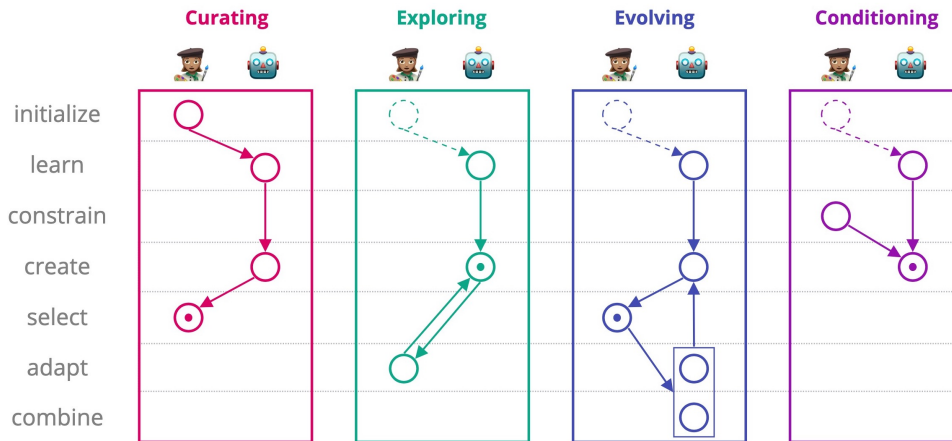


Figure 1: We present the four primary interaction patterns between human and GAN mapped along the vertical axis of actions. The circles display the taking of an action, with possible final actions depicted as points.

could support co-creative GAN applications. For example, Yildirim et al. [26] trained a GAN to generate dress designs based on given features such as color. Among all studies, we found four underlying interaction patterns for which we show the graphical notations in Figure 1. Interaction patterns are not mutually exclusive and can be combined. Table 1 provides an overview of how the surveyed examples apply the patterns.

3.2.1. Curating

The most straight-forward interaction pattern is categorized as Curating. After being *initialized* by a human designer, a GAN *learns* a to generate artifacts by being trained through gradient descent. Then, a set of outputs is *created* by sampling from the GAN, from which the human agent *selects* a set of final artifacts. Note that the selection can also include all generated artifacts. After the training, the GAN is the only one influencing the creation process. It does not receive any input from the human side. The main creative human input is given during the initialization. Hence, choosing a dataset and a model plays a crucial role in the human agent’s initiative.

Curating has been applied to produce new artifacts by adding variation to an existing dataset in several domains. For instance, fashion designers might purposely choose a collection of clothing designs as training distribution.¹ In that way, they determine the overall style of newly generated artifacts to which they aim to add novelty through the GAN’s variation. Kato et al. [32] train a GAN to produce outputs of similar quality to human-made designs.

¹<https://www.vogue.com/fashion-shows/fall-2020-menswear/acne-studios>

They replace the production step of design drawing in a fashion design workflow with GAN creations. Then, the output design is given to pattern makers.

Besides selecting the training data, the model architecture can also pose a design factor. Elgammal et al. [31] demonstrate that with art-generating *Creative Adversarial Networks* (CANs). They apply a style ambiguity loss when training the networks on an art dataset assigned to different art styles as categories. As a result, the CAN learns to output artifacts that deviate from existing art styles. While there are no limitations to the complexity of designing the initial GAN set-up, Curating is the simplest form for GANs to partake in the subsequent creative process.

3.2.2. Exploring

The interaction pattern Exploring describes an interaction where the human iteratively *adapts* artifacts that are created by sampling from the GAN. This is typically done by introducing alterations via the latent representations of artifacts. Altering the latent code of an artifact can be understood as providing a recipe to *create* a new artifact. By steering the following creation process, the designer guides the exploration path through the space of possible designs. The latent space can be explored by traversing along different directions. These directions cause different changes to the output features. Users can control the change in direction by e.g. moving a slider. We present three common directions to adapt artifacts below. Figure 2 provides a visual example per direction.

Interpolation. One direction to move an artifact through latent space is towards other artifacts in the




Figure 2: From left to right, we display examples of Exploring through *interpolation* between generated Zelda rooms [21], along *latent directions* of generated a Mario level [21], and along *semantic features* in CG-GAN [23].

space. In doing so, one explores intermediate alternatives between existing artifacts. Among other methods, Schrum et al. [21] apply interpolation to explore the design space of GANs trained to create tile-based rooms for the game *The Legend of Zelda*. By traversing between two different level designs, the game designer can prompt the GAN to create new designs lying in between. *CREA.blender* [33], a tool that tests users’ creativity in creating GAN outputs, allows the exploration of the intermediate space of up to five pre-selected images. Sliders let the user control the strength of an image’s representation in the generated result.

Latent directions. Instead of altering GAN-generated artifacts in relation to one another, they can also be adapted by directly increasing or decreasing their latent variables. As a result, the artifact moves along the GAN’s latent directions. In another part of Schrum et al.’s study, users adjust a game level design by moving a slider per latent variable. Because the latent variables do not correspond to visual output features, changes to the output design are not necessarily explainable, which users might experience as frustrating when having a target change in mind [21].

Semantic features. The adaption along directions in latent space that represent semantic attributes is better understandable. After finding such an attribute vector in the latent space, an artifact’s latent code can traverse along it. That influences the appearance of the attribute in the artifact. Research in the domain of semantic face editing [34, 35] has widely explored this technique to alter facial features. For example, the editing function in the interface of the *Composite Generating GAN* (CG-GAN) [23] supports interaction in that manner. Human agents adapt generated faces by adding or subtracting certain facial traits. They may choose to lock other facial features during that process to avoid correlated changes.

ArtBreeder [36], a website where users can make GAN-generated art, applies another approach to disentangle semantic features, resulting in a similar control to CG-GAN for a faces generator. By applying the *BigGAN* [37] or *StyleGAN* [7] architecture, categories present in the training data, such as facial attributes, can be disen-

gled during the training process. In practice, different attributes link to different layers and sliders let users manipulate target features. On top of that, Artbreeder allows users to introduce semantic directions based on sample images containing target features.

The features along which users explore the conceptual space can also be combined. For example, *CREA.blender SDG* [4], a tool to develop utopian and dystopian landscape images, lets users create interpolations between images as in *CREA.blender*. With the generator from a StyleGAN, the computational agent learned to distinguish between content and style of the environment. In the interaction, two sliders below each image let the human agent control how much of its content and style feed into the resulting design. Hence, the simple interpolation of artifacts expands to consider semantic properties.

Bau et al. [38] also bridge the gap between the internal representation of GANs and interpretable concepts. By analyzing how the latter encode in the former, they propose *GANDissect*, a method to identify latent units in the generator’s layers that activate certain aspects in the output image. Taken to practice in the application *GANPaint* [39], the authors propose an interface that lets users add, delete, or alter aspects in images. Using a brush to indicate target locations, or pixels, in the image, the human agent can adapt GAN-generated artifacts by indirectly de-/activating the corresponding latent units. By representing actual photographs as a GAN’s generation, the changes can be applied to existing images. This allows a human agent “to manipulate a photograph not with physical colors, but with abstract concepts such as object types and visual attributes” [39, p.2].

3.2.3. Evolving

In the third identified pattern, Evolving, the human *selects* artifacts for the computational agent to *adapt* and *combine* in the following step. Within the field of evolutionary computation, these actions are also referred to as mutation and recombination, respectively. This interaction utilizes interactive evolutionary computation [40], more specifically an interactive genetic algorithm (IGA). Here, the latent codes represent the genotype of artifacts,




Figure 3: From left to right, we display examples of Conditioning with a hand-drawn segmentation mask [42], a contour sketch [43], and an image and text description [25].

which are in turn considered as the phenotype. For a number of iterations, users make a selection from a population of artifacts to be adjusted for the next generation. In addition, users might be given the option to frame this process by specifying the exact action for an artifact or the operation parameters. While Exploring implies that the human agent directly controls the features that artifacts are adapted by, Evolving suggests they only select artifacts for the computational agent to develop further.

Bontrager et al. [41] let human agents evolve faces, shoes, and chairs in that manner. Users select artifacts from the first creation step. Their latent codes are then recombined through crossover and mutated by adding Gaussian noise. With the adapted latent codes, new items are sampled and the process repeats. With the power to guide the process through selecting artifacts to survive, the human agent optimizes the resulting artifacts according to their preferences. It is the computational agent that adjusts the artifacts primarily. However, the human agent can influence that, too, by specifying the rules for this traversal search [12]. In practice, this would imply setting algorithmic parameters such as mutation or crossover rate. Schrum et al. demonstrated that in the evolving functionality of their study [21]. Here, users select generated game levels to be recombined and mutated for the next generation, also referred to as selective breeding. By moving a slider, the user sets the rate for the applied polynomial mutation. In the evolving process underlying CG-GAN, the user can choose along which *semantic features* the algorithms should mutate artifacts. Similar directions for traversing the latent space as presented in section 3.2.2 could also be specified for the computational agent. The human agent might even select them during the design process. However, as the human agent does not directly adapt artifacts in this pattern, the directions do not concern the interaction and are not stated here.

3.2.4. Conditioning

Conditioning describes an interaction in which the human *constrains* the desired artifact characteristics before outputs are *created* through sampling from the GAN.

Hence, a requirement is that the GAN architecture can take conditions into account when sampling artifacts.

The constraint by the human can be expressed in several modalities, of which three examples are depicted in Figure 3 and described in the following. In the domain of fashion design, one can think of several parameters that might guide the creation of a clothing item. Yildirim et al. [42] for instance propose a GAN model that disentangles color, texture, and shape through different loss functions. During training, parts of the latent code are assigned to the three features. With this model, a user could pre-define the shape of a dress in a hand drawing. The drawing is then captured as a 512-dimensional mask and included into the input given to the GAN. Similarly, Xin and Arakawa [43] apply a GAN model that takes a contour image as input as part of their object design system. In that, the human designer can determine an outline before the GAN generates the artifact, such as the contour of a shoe. Zhao and Ma [44] note that providing a conditional GAN with a hand-drawn constraint can cause difficulties when the drawing style differs from the distribution of training drawings [44].

Setting landmarks of a drawing may help further define a target design. Therefore, the authors apply a second generative network in the constraint step, that transforms those landmarks into a compensation image. Together with the user’s drawing, it acts as a reference to guide the generation by the main GAN. The example shows that constraints can consist of multiple inputs, as is also the case in the following case for clothing design system suggested by Zhu et al. [25]. The authors let human designers guide the fashion design flow by giving an instruction in text form. First, a segmentation mask is derived from the photo of the person to be dressed. Together with the text encoding of an outfit description, the mask is given to their GAN model as input. This allows them to *dress* a person according to a human’s idea. The GAN contributes creative initiative through unpredictable variation.



Figure 4: From left to right, we display the interaction pattern for adding Conditioning to the flow of Xin & Arakawa’s design system [43], designs from a user study by Evolving alone, and designs generated by Evolving and Conditioning.

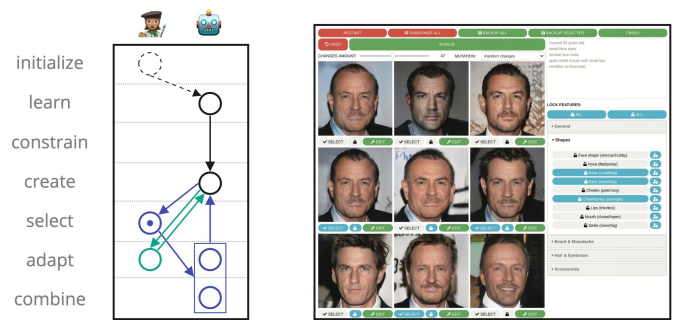


Figure 5: We display the interaction pattern of CG-GAN giving the user the possibility to step out of the Evolving loop for Exploring an image by clicking on the ‘EDIT’ button beneath one of the generated artifacts in the interface displayed on the right [23].

3.3. Combination of interaction patterns

The primary interaction patterns presented above are to be understood as an elementary set of interactions between human and GAN. As some of the reviewed applications show, they can be combined to create more advanced co-creative processes.

For instance, Xin and Arakawa [43] combine Evolving and Conditioning in their suggestion of an object design system. Figure 4 shows the interaction pattern for the system trained on datasets of shoes, handbags, and other fashion items. In the first step of their user study, users evolved shoe designs in an IGA without the option to constrain the design space, following simply the Evolving pattern. Secondly, they extend the generation process so that users could prompt the GAN’s creation with a contour image of the desired shoe (see Figure 3). By adding Conditioning to the Evolving interaction pattern, users can better derive designs that resemble the target sneaker, as the example in Figure 4 shows.

Instead of allowing the user to pre-limit the conceptual

space of a GAN, Zaltron et al. [23] give users the chance to adapt generated images. The interaction underlying their proposed CG-GAN combines the Evolving and Exploring pattern, as displayed in Figure 5. First, the model lets users evolve faces similar to Xin and Arakawa’s IGA for object design. But during the process, users can edit artifacts by stepping out of the evolutionary loop. In doing so, users fine-tune a generated design by traversing along semantic facial features with the help of sliders. The evolved artifact is then added back into the loop.

Oppositely, we can observe the Exploring pattern being expanded with Evolving functionalities in CREA.blender [33]. In its simplest form, CREA.blender asks users to develop animal-like shapes based on pre-selected basis images. Then, in the *open-play mode*, users can themselves select basis images by replacing the current ones. With the selected images as a starting point, they explore new images. In comparison to the “more goal-oriented creative tasks” [33, p.4], letting the user take part in the selection supports their “open-ended creativity” [33, p.4].

Table 1
Selection of GAN studies.

Application	Domain	Curating	Exploring	Evolving	Conditioning
Kato et al. [32]	clothing	•			
Elgammal et al. (CAN) [31]	art	•			
Bau et al. (GANPaint) [39]	photos		•		
Schrump et al. [21]	games		•	•	
Bontrager et al. (DeepIE) [41]	objects			•	
Xin and Arakawa [43]	objects			•	•
Zaltron et al. (CG-GAN) [23]	faces		•	•	
Simon (Artbreeder) [36]	diverse		•	•	
Gajdacz et al. (CREA.blender) [33]	objects		•	•	
Zhao and Ma [44]	clothing				•
Zhu et al. (FashionGAN) [25]	clothing				•
Yildirim et al. [42]	clothing				•

4. Discussion

In this section we discuss the co-creative affordance of the interaction patterns and reflect on the potential usage of the framework.

The interaction patterns we identified based on the framework allow us to analyze how co-creative GAN applications support co-creativity to different degrees. For example, where Muller et al. [14] identify one interaction pattern, we distinguish between Curating and Conditioning. More specifically, they categorize Elgammal et al.’s CAN and Zhao and Ma’s design method into the same interaction pattern characterized by the population of a solution space with a generative model. However, the two examples are fundamentally different in how they support co-creativity, and our proposed adaption allows us to make a fine-grained distinction between the two ways of interacting. Curating, as in the case of CAN, lets the human pre-determine the GAN’s design space through its initialization as an active design choice. Hence, the human agent’s primary input limits the GAN’s creative potential before training. Conditioning allows the human to narrow the existing design space of the GAN after training. We categorize Zhao and Ma’s example as such because the human agent steers the GANs by providing constraints, in this case, a drawing and landmarks, after the initial training. While Curating leaves the human agent out of the actual (co-)creative process and gives the GAN complete creative control, Conditioning restricts the GAN’s creativity by assigning the human agent more creative authority along the way, hence allowing for more co-creativity. By committing to a human-set frame, the GAN loses parts of its creative potential, leading to a trade-off between novelty and typicality with regards to humans’ expectations.

We regard Curating and Conditioning as *one-shot* approaches since the computational agent does not recall previously created artifacts, making each iteration an

isolated process. However, in the case of Conditioning, a user might map the constraint space and learn its *constraint language* by iterating the pattern repeatedly. While our preliminary framework does not currently capture this step, this points to the asymmetry in our framework, focusing on the GAN-specific technical grounding of actions rather than human-specific aspects.

Opposed to the one-shot patterns, Exploring and Evolving imply iterative interaction between human and GAN. While human initiative appears only after the first creation stage in both patterns, they support the co-creative exchange differently. Through selection in an Evolving flow, the human agent narrows the conceptual space of the GAN. Hence, they restrict its creative potential to developing a chosen population of designs. In comparison, Exploring leaves the conceptual space as is but lets the human designer traverse through it. However, the directions taken restrict the exploration. Freely altering the latent variables allows for a broader exploration than the interpolation between two artifacts, whereas traversing along semantic features limits the search to given paths. Overall, Exploring affords co-creativity by simply allowing the human agent to move in the conceptual space. In contrast, Evolving supports a more targeted search strategy for the user but restricts the GAN’s creative potential.

As the identified patterns support different purposes, their mapping allows for a reflection about their application and combination. When designing a process that starts from pre-defined characteristics, modeling the interaction according to the Conditioning pattern would let humans set the scope right from the beginning of a creation process. But when designing for a scenario that requires targeted search for distinguished properties in a design space, Evolving might be the pattern of choice. Exploring, on the other hand, could support a more inspirational excursion through artifact possibilities. When

combined however, the exploration of an evolved object as well as the evolution of previously conditioned objects lets the human agent fine-tune its properties.

Adapting Spoto and Oleynik's and Muller et al.'s frameworks allows us to make clear distinctions between the four prevailing interaction patterns. By tailoring the actions, we exchange the breadth of analyzing a wide range of MI-CC methods for the in-depth analysis of co-creative GAN applications. With the growing development of GANs for interaction, we hope that our framework supports the discovery of more GAN-specific interaction patterns. For example, a novel interaction pattern could be found in the recently suggested approach of *Rewriting GANs* [45], that allows users to update the generator's weights through interaction so the GAN learns during the process.

Awareness of the interaction patterns might inform the design of GAN applications according to the creative task at hand. Reflecting about them helps us facilitate meaningful interactions between GANs and humans that follow the purpose of a design system. The patterns might also point out alternative interactions by disclosing questions like: *How could a human agent participate in the direct creation of artifacts? Could agents (re-)learn during co-creation? What would an interaction pattern with more than two agents look like?* In practice, our framework may serve as a tool for UX designers to choose the correct interaction pattern when including GANs in creative design processes.

In addition, the framework enables constructive discussions on GANs' role in human-AI design collaborations as it allows us to follow how agents can restrict the creative potential of each other. Being able to trace the development of collaboratively created objects further allows us to understand where its properties originate. This insight matters when co-designing with intelligent systems trained on biased data distributions. As AI systems might further expose this bias through newly generated objects, reflecting on their participation is crucial for an informed creative process. While GANs become more apparent and easily applicable in creative tasks, this becomes essential.

5. Conclusion and Future Work

Our preliminary framework suggests mapping interactions as a basis for understanding co-creative GAN applications. We identified four primary interaction patterns, which we demonstrated along with examples of existing approaches. The patterns let us understand how different GAN applications support co-creativity. This insight can inform how we design co-creative GAN applications. The preliminary framework contributes to bridging the knowledge gap between machine learning engineers and

designers, who aim to apply GANs in creative processes. Additionally, it might serve as a starting point to discover how other features in GANs might offer novel interaction techniques. Future work in this direction could include more human-centered design steps.

References

- [1] R. T. Hughes, L. Zhu, T. Bednarz, Generative Adversarial Networks-Enabled Human-Artificial Intelligence Collaborative Applications for Creative and Design Industries: A Systematic Review of Current Approaches and Trends, *Frontiers in Artificial Intelligence* 4 (2021). doi:10.3389/frai.2021.604234, publisher: Frontiers Media SA.
- [2] T. Park, M.-Y. Liu, T.-C. Wang, J.-Y. Zhu, Semantic Image Synthesis with Spatially-Adaptive Normalization, in: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 2337-2346. URL: https://openaccess.thecvf.com/content_CVPR_2019/html/Park_Semantic_Image_Synthesis_With_Spatially-Adaptive_Normalization_CVPR_2019_paper.html.
- [3] A. Noyman, K. Larson, DeepScope: HCI Platform for Generative Cityscape Visualization, in: *Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems, CHI EA '20*, Association for Computing Machinery, New York, NY, USA, 2020, pp. 1-9. URL: <https://doi.org/10.1145/3334480.3382809>. doi:10.1145/3334480.3382809.
- [4] J. Rafner, S. Langsfjord, A. Hjorth, M. Gajdacz, L. Philipsen, S. Risi, J. Simon, J. Sherson, Utopian or Dystopian?: using a ML-assisted image generation game to empower the general public to envision the future, in: *Creativity and Cognition*, ACM, Virtual Event Italy, 2021, pp. 1-5. URL: <https://dl.acm.org/doi/10.1145/3450741.3466815>. doi:10.1145/3450741.3466815.
- [5] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative Adversarial Nets, in: *NIPS'14: Proceedings of the 27th International Conference on Neural Information Processing Systems*, volume 2, 2014, pp. 2672-2680. URL: <http://www.github.com/goodfeli/adversarial>.
- [6] A. Radford, L. Metz, S. Chintala, Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks (2015). URL: <http://arxiv.org/abs/1511.06434>.
- [7] T. Karras, S. Laine, T. Aila, A Style-Based Generator Architecture for Generative Adversarial Networks, in: *Proceedings of the IEEE/CVF conference on*

- computer vision and pattern recognition, 2019, pp. 4401–4410. URL: <http://arxiv.org/abs/1812.04948>.
- [8] J. Zhu, A. Liapis, S. Risi, R. Bidarra, G. M. Youngblood, Explainable AI for Designers: A Human-Centered Perspective on Mixed-Initiative Co-Creation, in: 2018 IEEE Conference on Computational Intelligence and Games (CIG), IEEE, Maastricht, 2018, pp. 1–8. URL: <https://ieeexplore.ieee.org/document/8490433/>. doi:10.1109/CIG.2018.8490433.
- [9] Q. Yang, A. Steinfeld, C. Rosé, J. Zimmerman, Re-examining Whether, Why, and How Human-AI Interaction Is Uniquely Difficult to Design, in: Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems, CHI '20, Association for Computing Machinery, New York, NY, USA, 2020, pp. 1–13. URL: <https://doi.org/10.1145/3313831.3376301>. doi:10.1145/3313831.3376301.
- [10] D. Buschek, L. Mecke, F. Lehmann, H. Dang, Nine Potential Pitfalls when Designing Human-AI Co-Creative Systems, in: Joint Proceedings of the ACM IUI 2021 Workshops, College Station, USA, 2021. URL: <http://arxiv.org/abs/2104.00358>.
- [11] A. Liapis, G. Smith, N. Shaker, Mixed-initiative content creation, in: Procedural Content Generation in Games, Springer International Publishing, Cham, 2016, pp. 195–214. URL: http://link.springer.com/10.1007/978-3-319-42716-4_11. doi:10.1007/978-3-319-42716-4_11, series Title: Computational Synthesis and Creative Systems.
- [12] G. N. Yannakakis, A. Liapis, C. Alexopoulos, Mixed-Initiative Co-Creativity, Fort Lauderdale, 2014, pp. 1–8.
- [13] A. Spoto, N. Oleynik, Library of Mixed-Initiative Creative Interfaces, 2017. URL: <http://mici.codingconduct.cc/>.
- [14] M. Muller, J. D. Weisz, W. Geyer, Mixed Initiative Generative AI Interfaces: An Analytic Framework for Generative AI Applications, in: Proceedings of the Workshop The Future of Co-Creative Systems - A Workshop on Human-Computer Co-Creativity of the 11th International Conference on Computational Creativity (ICCC 2020), 2020. URL: https://computationalcreativity.net/workshops/cocreative-iccc20/papers/Future_of_co-creative_systems_185.pdf.
- [15] J. Zhu, J. Villareale, N. Javvaji, S. Risi, M. Löwe, R. Weigelt, C. Harteveld, Player-AI Interaction: What Neural Network Games Reveal About AI as Play, in: Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems, 2021, pp. 1–17. URL: <https://doi.org/10.1145/3411764.3445307>, arXiv: 2101.06220.
- [16] X. Wu, K. Xu, P. Hall, A survey of image synthesis and editing with generative adversarial networks, Tsinghua Science and Technology 22 (2017) 660–674. doi:10.23919/TST.2017.8195348, conference Name: Tsinghua Science and Technology.
- [17] P. Shamsolmoali, M. Zareapoor, E. Granger, H. Zhou, R. Wang, M. E. Celebi, J. Yang, Image synthesis with adversarial networks: A comprehensive survey and case studies, Information Fusion 72 (2021) 126–146. URL: <https://www.sciencedirect.com/science/article/pii/S1566253521000385>. doi:10.1016/j.inffus.2021.02.014.
- [18] W. Roy, G. Kelly, R. Leer, F. Ricardo, A Survey on Adversarial Image Synthesis, ACM Comput. Surv. (2021). URL: <http://arxiv.org/abs/2106.16056>, arXiv: 2106.16056.
- [19] Q. Zhang, S.-C. Zhu, Visual Interpretability for Deep Learning: a Survey, Frontiers of Information Technology & Electronic Engineering volume (2018). URL: <http://arxiv.org/abs/1802.00614>, arXiv: 1802.00614.
- [20] S. Shahriar, GAN Computers Generate Arts? A Survey on Visual Arts, Music, and Literary Text Generation using Generative Adversarial Network, arXiv:2108.03857 [cs, stat] (2021). URL: <http://arxiv.org/abs/2108.03857>, arXiv: 2108.03857.
- [21] J. Schrum, J. Gutierrez, V. Volz, J. Liu, S. Lucas, S. Risi, Interactive Evolution and Exploration Within Latent Level-Design Space of Generative Adversarial Networks, in: Proceedings of the 2020 Genetic and Evolutionary Computation Conference, 2020, pp. 148–156. URL: <https://doi.org/10.1145/3377930.3389821>, arXiv: 2004.00151.
- [22] L. Sun, P. Chen, W. Xiang, P. Chen, W.-y. Gao, K.-j. Zhang, SmartPaint: a co-creative drawing system based on generative adversarial networks, Frontiers of Information Technology & Electronic Engineering 20 (2019) 1644–1656. URL: <http://link.springer.com/10.1631/FITEE.1900386>. doi:10.1631/FITEE.1900386.
- [23] N. Zaltron, L. Zurlo, S. Risi, CG-GAN: An Interactive Evolutionary GAN-Based Approach for Facial Composite Generation, Proceedings of the AAAI Conference on Artificial Intelligence 34 (2020) 2544–2551. URL: <https://aaai.org/ojs/index.php/AAAI/article/view/5637>. doi:10.1609/aaai.v34i03.5637.
- [24] D. Friedman, D. Pollak, Image Co-Creation by Non-Programmers and Generative Adversarial Networks, in: Joint Proceedings of the ACM IUI 2021 Workshops, 2021, p. 10.
- [25] S. Zhu, S. Fidler, R. Urtasun, D. Lin, C. C. Loy, Be Your Own Prada: Fashion Synthesis with Structural Coherence, in: Proceedings of the IEEE International Conference on Computer Vision (ICCV), 2017, pp. 1680–1688. URL: <http://arxiv.org/abs/1710.07346>, arXiv: 1710.07346.

- [26] G. Yildirim, N. Jetchev, R. Vollgraf, U. Bergmann, Generating High-Resolution Fashion Model Images Wearing Custom Outfits, 2019. URL: <http://arxiv.org/abs/1908.08847>.
- [27] B. Roziere, F. Teytaud, V. Hosu, H. Lin, J. Rapin, M. Zameshina, O. Teytaud, EvolGAN: Evolutionary Generative Adversarial Networks, 2020. URL: <http://arxiv.org/abs/2009.13311>.
- [28] Z. Zeng, X. Sun, X. Liao, Artificial Intelligence Augments Design Creativity: A Typeface Family Design Experiment | SpringerLink, in: Design, User Experience, and Usability. User Experience in Advanced Technological Environments, 2019. URL: https://link.springer.com/chapter/10.1007/978-3-030-23541-3_29.
- [29] S. Deterding, J. Hook, R. Fiebrink, M. Gillies, J. Gow, M. Akten, G. Smith, A. Liapis, K. Compton, Mixed-Initiative Creative Interfaces, in: Proceedings of the 2017 CHI Conference Extended Abstracts on Human Factors in Computing Systems, ACM, Denver Colorado USA, 2017, pp. 628–635. URL: <https://dl.acm.org/doi/10.1145/3027063.3027072>. doi:10.1145/3027063.3027072.
- [30] A. Liapis, G. N. Yannakakis, Boosting Computational Creativity with Human Interaction in Mixed-Initiative Co-Creation Tasks, in: Proceedings of the ICCG workshop on Computational Creativity and Games, Paris, 2016.
- [31] A. Elgammal, B. Liu, M. Elhoseiny, M. Mazzone, CAN: Creative Adversarial Networks, Generating "Art" by Learning About Styles and Deviating from Style Norms (2017). URL: <http://arxiv.org/abs/1706.07068>.
- [32] N. Kato, H. Ozone, K. Oomori, C. W. Ooi, Y. Ochiai, GANs-based Clothes Design: Pattern Maker Is All You Need to Design Clothing, in: Proceedings of the 10th Augmented Human International Conference 2019, ACM, Reims France, 2019, pp. 1–7. URL: <https://dl.acm.org/doi/10.1145/3311823.3311863>. doi:10.1145/3311823.3311863.
- [33] M. Gajdacz, J. Rafner, S. Langsford, A. Hjorth, C. Bergenholtz, M. M. Biskjaer, L. Noy, S. Risi, J. Sherson, CREA.blender: a GAN based casual creator for creativity assessment, in: Proceedings of the International Conference on Computational Creativity (ICCC '21), 2021, p. 5.
- [34] E. Denton, B. Hutchinson, M. Mitchell, T. Geburu, A. Zaldivar, Image Counterfactual Sensitivity Analysis for Detecting Unintended Bias (2019). URL: <http://arxiv.org/abs/1906.06439>.
- [35] Y. Shen, J. Gu, X. Tang, B. Zhou, Interpreting the Latent Space of GANs for Semantic Face Editing, in: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), IEEE, Seattle, WA, USA, 2020, pp. 9240–9249. URL: <https://ieeexplore.ieee.org/document/9157070/>. doi:10.1109/CVPR42600.2020.00926.
- [36] J. Simon, Artbreeder, 2022. URL: <https://www.artbreeder.com/browse>.
- [37] A. Brock, J. Donahue, K. Simonyan, Large Scale GAN Training for High Fidelity Natural Image Synthesis, arXiv:1809.11096 [cs, stat] (2019). URL: <http://arxiv.org/abs/1809.11096>, arXiv: 1809.11096.
- [38] D. Bau, J.-Y. Zhu, H. Strobelt, B. Zhou, J. B. Tenenbaum, W. T. Freeman, A. Torralba, GAN Dissection: Visualizing and Understanding Generative Adversarial Networks, in: Proceedings of the International Conference on Learning Representations (ICLR), 2019, p. 19.
- [39] D. Bau, H. Strobelt, W. Peebles, J. Wulff, B. Zhou, J.-Y. Zhu, A. Torralba, Semantic photo manipulation with a generative image prior, ACM Transactions on Graphics 38 (2019) 1–11. URL: <https://dl.acm.org/doi/10.1145/3306346.3323023>. doi:10.1145/3306346.3323023.
- [40] H. Takagi, Interactive evolutionary computation: fusion of the capabilities of EC optimization and human evaluation, Proceedings of the IEEE 89 (2001) 1275–1296. URL: <http://ieeexplore.ieee.org/document/949485/>. doi:10.1109/5.949485.
- [41] P. Bontrager, W. Lin, J. Togelius, S. Risi, Deep Interactive Evolution, in: International Conference on Computational Intelligence in Music, Sound, Art and Design : EvoMUSART 2018, 2018.
- [42] G. Yildirim, C. Seward, U. Bergmann, Disentangling Multiple Conditional Inputs in GANs, in: KDD 2018 Conference AI for Fashion Workshop, 2018. URL: <http://arxiv.org/abs/1806.07819>, arXiv: 1806.07819.
- [43] C. Xin, K. Arakawa, Object Design System by Interactive Evolutionary Computation Using GAN with Contour Images, in: Human Centred Intelligent Systems - Proceedings of KES-HCIS 2021 Conference, volume 244, Springer Singapore, Singapore, 2021, pp. 66–75. URL: https://link.springer.com/10.1007/978-981-16-3264-8_7. doi:10.1007/978-981-16-3264-8_7, series Title: Smart Innovation, Systems and Technologies.
- [44] Z. Zhao, X. Ma, A Compensation Method of Two-Stage Image Generation for Human-AI Collaborated In-Situ Fashion Design in Augmented Reality Environment, in: 2018 IEEE International Conference on Artificial Intelligence and Virtual Reality (AIVR), IEEE, Taichung, Taiwan, 2018, pp. 76–83. URL: <https://ieeexplore.ieee.org/document/8613637/>. doi:10.1109/AIVR.2018.00018.
- [45] D. Bau, S. Liu, T. Wang, J.-Y. Zhu, A. Torralba, Rewriting a Deep Generative Model, in: Proceedings of the European Conference on Computer Vision (ECCV), 2020.