

Counterfactual Edits for Generative Evaluation

Maria Lymperaïou^{1,*}, Giorgos Filandrianos^{1,*}, Konstantinos Thomas^{1,*} and Giorgos Stamou¹

¹*AILS Laboratory, School of Electrical and Computer Engineering, National Technical University of Athens*

Abstract

Evaluation of generative models has been an underrepresented field despite the surge of generative architectures. Most recent models are evaluated upon rather obsolete metrics which suffer from robustness issues, while being unable to assess more aspects of visual quality, such as compositionality and logic of synthesis. At the same time, the explainability of generative models remains a limited, though important, research direction with several current attempts requiring access to the inner functionalities of generative models. Contrary to prior literature, we view generative models as a black box, and we propose a framework for the evaluation and explanation of synthesized results based on concepts instead of pixels. Our framework exploits knowledge-based counterfactual edits that underline which objects or attributes should be inserted, removed, or replaced from generated images to approach their ground truth conditioning. Moreover, global explanations produced by accumulating local edits can also reveal what concepts a model cannot generate in total. The application of our framework on various models designed for the challenging tasks of Story Visualization and Scene Synthesis verifies the power of our approach in the model-agnostic setting.

Keywords

Image Generation, Counterfactual Explanations, Diffusion Models, Story Visualization, Generative Evaluation, XAI

1. Introduction

Image generation has been one of the most popular deep learning tasks, inspiring many impressive state-of-the-art applications [1, 2, 3, 4, 5, 6, 7]. Even since the introduction of Generative Adversarial Networks (GANs) [8], which marked one of the first significant breakthroughs in the field, most applications focused on enhancing image quality according to human perception. At the same time, the automatic evaluation of the generated samples remains a long-standing problem as there are no ground truth data to measure against. The valuation of such generative tasks, so far, relies on pixel-level metrics such as Inception Score (IS) [9], Frechet Inception Distance (FID) [10], Learned Perceptual Image Patch Similarity (LPIPS) [11], to provide a quality measure for the generated samples. Consequently, the list of literature evaluated upon those benchmark metrics is long; yet concerns have been raised that their brittleness [12] is leading

In A. Martin, K. Hinkelmann, H.-G. Fill, A. Gerber, D. Lenat, R. Stolle, F. van Harmelen (Eds.), Proceedings of the AAAI 2023 Spring Symposium on Challenges Requiring the Combination of Machine Learning and Knowledge Engineering (AAAI-MAKE 2023), Hyatt Regency, San Francisco Airport, California, USA, March 27-29, 2023.


*Corresponding author.

✉ marialymp@islab.ntua.gr (M. Lymperaïou); geofila@islab.ntua.gr (G. Filandrianos); kthomas@islab.ntua.gr (K. Thomas); gstam@cs.ntua.gr (G. Stamou)

ORCID 0000-0001-9442-4186 (M. Lymperaïou); 0000-0002-7489-7776 (K. Thomas); 0000-0003-1210-9874 (G. Stamou)



© 2023 Copyright 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)

to inaccurate results. Although recent metrics, such as Clean-FID [12] can resolve some issues regarding visual artifacts, they still cannot address major issues such as the evaluation of complex images, compositionality, logic, and fairness of generation [13]. Moreover, when it comes to conditional generation, we further require a measure of whether objects and attributes mentioned in the conditioning are successfully depicted on the generated samples. Current attempts in conditional synthesis evaluation remain limited [14, 15] while still facing the shortcomings of their unconditional counterparts, on which they are built.

Explainability of generative models is another emerging field, which has currently received way less attention compared to discriminative models [16, 17]. The incorporation of explainable feedback in Generative Adversarial Networks (GANs) has demonstrated a promising research direction [18], while other works focus on interpreting GAN inner structure [19]. Overfitting in GANs can be tackled by determining the image areas that contributed to classifying a sample as fake/real, thus explaining the discriminator’s decision [20]. This limited literature impedes the development of *explainable evaluation* for generative models, even though related attempts have gained ground in other AI domains, such as Natural Language Processing [21, 22, 23].

We argue that resolving generative evaluation challenges calls for a conceptual approach to the evaluation process, diverging from the pixel-level route. Relying on concepts instead of pixels offers the advantage of enhanced interpretability regarding the evaluation process and paves the way for *explainable evaluation* of generative models. Identifying concepts (objects or attributes) that can or cannot be generated reveals the capabilities and biases of the model at hand, thus driving potential architectural modifications. In this paper, we present the first *explainable evaluation* technique targeting generative models. Specifically, we utilize *counterfactual explanations* to frame conditional generative evaluation as the answer to the following question: *What concepts need to change in a generated sample X, for it to reach its conditioning c?* Conceptual edits guided from external knowledge sources [24] have shown to efficiently indicate the shortest possible path to reach the conditioning concepts. Furthermore, existing works that combine explainability with image generation operate on specific models [18, 19, 20] and demand access to their inner structure (white-box techniques), while our approach only requires generated outputs along with their ground truth conditioning, yet still regarding the generative model as a black-box. We, therefore, contribute to the following:

1. We propose the first *conceptual* rather than pixel-based generative evaluation framework¹, suitable for various tasks such as Scene Generation (SG) and Story Visualization (SV).
2. Our metrics are *explainable* by design, illustrating which concepts need to be inserted, deleted, or replaced in the generated images, for them to approach the ground truth conditioning. Those edit operations are applied in a model-agnostic setting, totally trespassing any access to the generative model inner workings.
3. Global explanations automatically reveal possible *blind spots* of generative models, i.e. concepts that a model is intrinsically incapable of generating.

2. Related work

Generative Adversarial Networks (GANs) [8] consist of two neural networks, a generator

¹<https://github.com/geofila/Counterfactual-Edits-for-Generative-Evaluation>

$G(z; \theta_g)$ and a discriminator $D(x; \theta_d)$. G maps random noise z , generated from a prior distribution $z \sim p_z$, to the data space. D , on the other hand, maps a sample x_i from the same data space to a scalar value $p_i = D(x_i)$, which represents the probability that the sample was drawn from the real data distribution. In the case of conditional GANs (cGANs) [25], G is fed not only with random noise z , but also with an additional conditioning vector y , which helps guide the generation of samples from specific sub-regions of the target distribution.

Several image generation cGANs [26, 27] perform well when it comes to generating images with distinct textures and colors. However, they tend to struggle with generating coherent overall object structures and other long-range dependencies, due to the limited nature of convolutional filters. The Self-Attention GAN (SAGAN) [28] was proposed as a solution to this problem; it utilizes a self-attention module in both G and D , as well as modern stabilization techniques such as Spectral Normalization of weights [29], while it leverages the two-timescale update rule [30] to impose different learning rates for G and D .

Diffusion models are breaking new ground in the field of conditional image generation and are becoming the state-of-the-art in that area [2]. These models work by adding noise to an image and then learning to reconstruct it. In the past year, there have been several exciting developments in the field of diffusion-based image synthesis. Stable Diffusion [31] allows for high-quality image synthesis even under resource constraints by applying the diffusion process in the latent space of autoencoders instead of at the pixel level in the image space. DALL-E2 [4] builds upon the success of its predecessor [3] by incorporating text-conditioned image embeddings learned from CLIP [32] as input to a diffusion model that acts as the decoder. The resulting images are photorealistic and accurate representations of the input text, and it also allows for language-guided manipulation of a source image. The work of Imagen [33] leverages large pre-trained language models, such as T5 [34], for language encoding and conducts image synthesis using the diffusion process. DreamBooth [35] takes Imagen a step further by allowing for context-aware image synthesis, given a text description of the desired context. This allows for the generation of various visual subjects while maintaining high image synthesis quality.

Conditional image synthesis has come a long way since the early days of text-conditioned image generation [36, 37]. First attempts produced images lacking in detail and quality. StackGAN [38] was the first model to significantly improve the quality of generated images using a multi-stage adversarial training process, followed by StackGAN++ [39] which further enhanced generation results. AttnGAN [40] employed attention mechanisms to generate fine-grained details in images based on individual words in the input text. SEGAN [41] took this a step further by only focusing attention on relevant keywords in the input text. DM-GAN [42] improved the quality of generated images by addressing fuzzy areas.

StoryGAN [43] is a generative model that synthesizes images based on sequential input (Story Visualization), using an RNN structure to encode the input text and provide context information to the generator. The generator is trained adversarially against two discriminators: the image discriminator, which evaluates image quality and text-image relevance, and the story discriminator, which ensures consistency across images given the entire story context. Recent work has focused on improving the baseline StoryGAN model [44] and exploring alternative

story encoding methods, such as using Transformer architectures [45, 46, 47].

Generative evaluation Despite the rapid advancements in image synthesis, generative evaluation is falling behind due to outdated evaluation practices [9, 10, 11, 48], mainly followed for benchmarking reasons, ignoring the problems recognized in recent works [12, 13]. Explainability in generative modeling can deliver interesting insights, though current efforts either remain model-specific [49, 18, 19, 20] or require discovering interpretable latent directions [50, 51, 52, 53], which is a non-trivial task. Our method serves both the evaluation and explainability of generative models under a single framework and is capable of being adapted to any generative model - even the ones serving sequential image generation [43, 45, 46, 47] - as it focuses solely on input and output concept sets.

Counterfactual explanations Contemporary AI research moves towards explaining a neural network’s train of ‘thought’, thus eXplainable AI (XAI) becomes a field of increasing interest [54]. Counterfactual explanations provide alternative realities based on minimal input modifications, hence revealing reasoning paths. Generative models are a straightforward approach when visual alternatives are explored [55]. Any alteration should be feasible with respect to original data distribution, an observation that adds constraints in alternative inputs [56]. Minimum alterations can be decided either by interfering with the black-box nature of neural networks [57], or not [24]. We chose to follow the black-box route, using counterfactual explanations to uncover clues on the reasoning processes of generative models.

3. Conceptual edits as counterfactual explanations

Our overall approach is heavily inspired by Filandrianos et al. [24], which explores the fundamental question of counterfactual reasoning: “What is the *minimal* change that has to occur in order for an image I to be classified as X instead of Y ?”, where X and Y are predicted categories of a pre-defined image classifier F . In our case, F is not necessary, since we by default place all generated concepts s in a set S , and all ground truth concepts t in a set T . Counterfactual explanations are capable of addressing the aforementioned question, providing the minimum number of conceptual edits to achieve the $S \rightarrow T$ transition for all $s \in S, t \in T$.

Concept distances instruct the shortest path that connects two specific concepts. Concept hierarchies are employed, deterministically defining the transition cost between concepts. We explore both the option to use external hierarchical knowledge such as WordNet [58], mapping extracted concepts to synsets, or alternatively handcraft specific hierarchies to allow highly controlled semantic distance definition. In both cases, we denote as $d(s, t)$ the distance between concepts s and t . There are three available concept **edit operations** to realize transitions:

- **Replacement (R)** $e_{s \rightarrow t}(S)$: A concept $s \in S$ is replaced with a concept $t \notin S$.
- **Deletion (D)** $e_{s-}(S)$: A concept $s \in S$ is deleted from S .
- **Insertion (I)** $e_{s+}(S)$: A concept $s \in S$ is inserted in S .

Each edit operation inherits the concept distances imposed by the selected hierarchy. Therefore, **R** operation considers the path between s and t so that $\min(d(s, t))$ is ensured. As in

[24], d also ensures *actionability* of edits, allowing semantically meaningful transitions (e.g. 'food' → 'pasta'), while prohibiting meaningless ones (e.g. 'food' → 'sky'). **D** and **I** operations regard the *root node* of the hierarchy as t and s respectively; in the case of WordNet, entity.n.01 serves as the root. **Concept Set Edit Distance (CSED)** $D(S \rightarrow T)$ is obtained by aggregating all possible minimum cost edit operations so that $S \rightarrow T$ is finally achieved:

$$CSED = D(S \rightarrow T) = \min \sum_{s \neq t}^{S, T} \sum_{R, D, I} d(s, t) \quad (1)$$

4. Method

The heart of our method consists of a pre-trained black-box generative model M which receives a semantic description c (in natural language or in symbolic format) as conditioning and produces an image I corresponding to c . We then use off-the-self automatic methods such as object detection, semantic segmentation, and others, in order to extract all the concepts depicted in the generated images I and append them in the *generated* (or *source*) concept set S . Similarly, concepts extracted from c contribute to the *real* (or *target*) concept set T . The format of c defines the concept extraction technique that is followed, ranging from linguistic concept extraction, if c is a textual sentence, to simple preprocessing, if c is already in a set format. Ultimately, we aspire to answer the following: "What are the *minimal* required changes in order to traverse from S to T ?" The outline of our method is presented in Figure 1.

4.1. Generative evaluation

The counterfactual backbone described in Section 3 highlights our process for generative evaluation, which we employ on two difficult tasks of the generative literature: **Story Visualization (SV)** and **Scene Generation (SG)**.

Story Visualization (SV) targets the sequential creation of images I_1, I_2, \dots, I_L that correspond, one-to-one, to a given sequential conditioning c_1, c_2, \dots, c_L of a total length L . The generated images need not only to remain *faithful* to their conditioning, but to also maintain serial *consistency*. We therefore define the two desiderata applicable to SV:

- **Faithfulness:** objects and attributes mentioned in c_k should also appear in frame I_k , for any story frame k , where $1 \leq k \leq L$.
- **Consistency:** objects or attributes appearing in frame I_k cannot disappear or change in later frames I_{k+1}, \dots, I_L , for any story frame k , where $1 \leq k < L$.

Since well-defined semantics are tied to counterfactual explanations [59], we regard CLEVR-SV [60] as the ideal dataset to demonstrate our approach, as it provides a set of concepts \mathcal{C} : shape (*cube, sphere, cylinder*), size (*small, large*), material (*rubber, metal*) and one of 8 colors (*blue, cyan, brown, yellow, red, green, purple, gray*). Each CLEVR-SV object contains $|\mathcal{C}|=4$ concepts

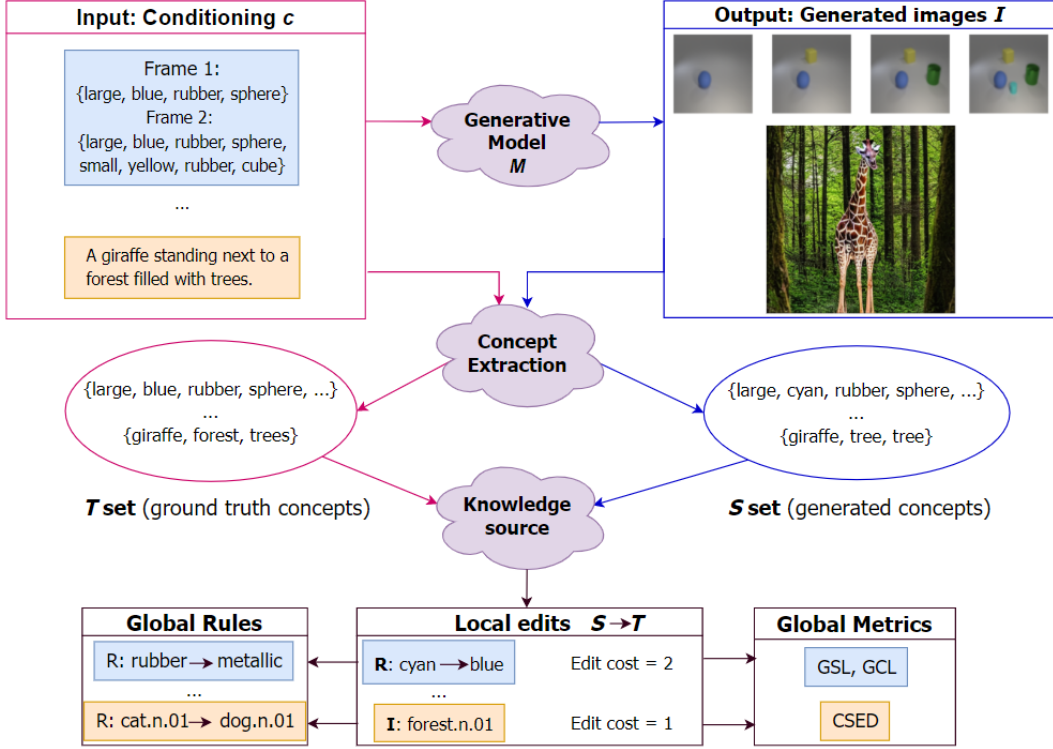


Figure 1: Outline of the proposed framework.

that describe its shape, size, material and color. We handcraft a simple hierarchy to group object semantics to generic concept classes, demonstrating the following inclusion relationships:

$$\begin{aligned}
 & (large, small) \subset Size \\
 & (blue, yellow, brown, grey, green, purple, cyan, red) \subset Color \\
 & (metallic, rubber) \subset Material \\
 & (sphere, cube, cylinder) \subset Shape
 \end{aligned} \tag{2}$$

CLEVR-SV contains stories of length $L=4$, with the k -th frame strictly containing k objects. Any of the three available edit operations can be relevant per frame: **D** of a concept, when a generated frame contains more objects than its ground truth match; **I** of a concept in the opposite case; **R** equals to a **D** followed by an **I**, and can be applied on frames with proper number of objects when semantics differ. In the default case, we assign equal costs of 1 for all semantics, as well as for **D** and **I** operations (**R** cost is the sum of **D** and **I** costs).

To measure story **faithfulness** we propose the **Story Loss (SL)** metric, which sums up the per-frame Concept Set Edit Distance ($CSED_k$) for $k = 1, 2, \dots, L$ frames of the story. Generated CLEVR-SV semantics for shape, size, material and color for the k -th frame form the concepts set S_k , while the semantics of the conditioning form T_k , with their $CSED_k$ denoted as $D(S_k, T_k)$. Thus, the cost for the transition $\{S_1, S_2, S_3, S_4\} \rightarrow \{T_1, T_2, T_3, T_4\}$ corresponding to the minimum cost **R D I** edits needed to transform the semantics of the generated sequence

$\{I_1, I_2, I_3, I_4\}$ to the semantics of its conditioning $\{c_1, c_2, c_3, c_4\}$ can be expressed as:

$$SL = \sum_{k=0}^{k=L} CSED_k = \sum_{k=0}^{k=L} D(S_k, T_k), \quad L = 4 \quad (3)$$

By scaling up the calculation of SL for a dataset containing N stories, we obtain the **Global Story Loss (GSL)** metric:

$$GSL = \sum_{i=0}^N SL_i = \sum_{i=0}^N \sum_{k=1}^{k=L} D(S_k, T_k)_i \quad (4)$$

As for story **consistency**, we propose the metric of **Consistency Loss (CL)**: the frame I_k is compared with I_{k-1} , $k = 2, \dots, L$ frames of the generated sequences to capture changes of semantics. A challenging aspect of CL is that there does not exist a ground truth concept set. However, since it is known by task definition that the k -th frame contains k objects, and the cardinality $|\mathcal{C}|$ of dataset concepts is predefined ($|\mathcal{C}|=4$ in the case of CLEVR-SV), we can assume that every previous frame constitutes the 'ground truth' corresponding to the concept set T . Commencing from the $k=1$ frame, we expect the cardinality of T to be equal with $|\mathcal{C}| \cdot k = |\mathcal{C}|$. Any discrepancy results in a penalty $p_k = |T| - |\mathcal{C}| \cdot k = CL_k$ for $k=1$. For later frames, we define as S the concept set corresponding to the k -th frame, and as T the 'ground truth' set comprised of the $k-1$ frame concepts. Mathematically, CL can be written as:

$$CL = p_{k=1} + \sum_{k=2}^{k=L} D(S_k, T_k), \quad T_k = S_{k-1}, L = 4 \quad (5)$$

In the ideal case, when the k -th frame contains k objects with \mathcal{C} semantics, we expect that $p_{k=1} = 0$ and $CL_{k>1} = |\mathcal{C}| \cdot (k - 1)$. By extending CL to N stories, **Global Consistency Loss (GCL)** evaluates the consistency capabilities of a generative model M in total:

$$GCL = \sum_{i=0}^N CL_i = \sum_{i=0}^N \{p_{k=1,i} + \sum_{k=2}^{k=L} D(S_k, T_k)_i, \quad T_{k,i} = S_{k-1,i}\} \quad (6)$$

Average values can be obtained for both local (**SL/CL**) and global (**GSL/GCL**) metrics:

$$\text{Avg } SL = \frac{1}{k} SL, \quad \text{Avg } GSL = \frac{1}{N} [\text{Avg } SL] = \frac{1}{N} GSL \quad (7)$$

For consistency, instead of exporting an average value over $\sum CL_k$, it is more meaningful to count how many times the $p_{k=1} = 0$, $CL_{k>1} = |\mathcal{C}| \cdot (k - 1)$ requirement was not respected, averaged for $L = k$ frames:

$$\text{Avg } CL = \frac{p_{k=1}}{k} + \frac{1}{k} \sum_{k=1}^{k=L} [CL_{k>1} \neq |\mathcal{C}| \cdot (k - 1)], \quad \text{Avg } GCL = \frac{1}{N} [\text{Avg } CL] \quad (8)$$

SL and **CL** are by nature *explainable*, as they do not only provide a measure of quality but also reveal the $S_k \rightarrow T_k$ edit paths. Those paths serve as *local counterfactual explanations*, highlighting

the erroneously generated semantics for this particular story, either in terms of **faithfulness** or **consistency**. Overall, higher **SL/GSL** and **CL/GCL** values denote lower conceptual generation quality. **GSL/GCL** edit paths correspond to *global counterfactual explanations*: rule extraction techniques provide frequent patterns, summarizing the behavior of M under investigation. Frequent **GSL** edit paths in fact contain *common misconceptions*, i.e. conditioning concepts that M cannot easily generate. Similarly, **GCL** edit paths reveal frequent *inconsistency patterns*, showcasing concepts that arbitrarily change within the story frames. Hence, by researching the question "What has to *minimally* change in order to transit from S to T ?", we eventually answer a more generic one: "Which concepts cannot be generated or preserved by M ?"

Scene Generation (SG) aims to synthesize a visual scene I based on a conditioning c . The synthesized image comprises multiple objects which interact with each other. Scene objects are also accompanied by attributes. The given conditioning c is more complex compared to conditionings provided for SV, since the concepts to be generated are numerous and not predefined; this yields a concept set \mathcal{C} of unknown but comparatively large cardinality.

COCO dataset [61] provides the ideal setting for evaluating generative **faithfulness** for SG, providing textual captions c that can serve as conditioning. We focus our endeavors on state-of-the-art open source diffusion models [2] from Huggingface², and specifically on Stable Diffusion v1.4 & v2 [62, 63] and Protogen x3.4 & 5.8 [64, 65] (details in Appendix). These models produce *realistic* images - an important aspect of the concept extraction (object detection) stage. We omit older SG architectures [66, 67, 68, 69, 70, inter alia] due to their inferior visual quality and their reliance on scene graphs and layouts for ensuring proper composition.

In the concept extraction stage, YOLO-v8 [71] and YOLOs [72] object detectors are leveraged to construct the generated concept set S . Since c is in textual format, spaCy [73] is used to extract *ground truth concepts* from captions that form the target concept set T . The semantically complex nature of concept distances related to COCO concepts requires a rich knowledge scheme, such as WordNet. For example, if c refers to concepts such as 'food' or 'animal', a diffusion model may generate more refined 'food' or 'animal' instances, for example, 'pasta' and 'dog' respectively. The object detectors will then return these refined classes, inducing some noise in the transformation process. Hierarchical knowledge can eliminate such issues: even though $T = \{food, animal\} \neq S = \{pasta, dog\}$, the two sets are **semantically equivalent** if we consider the hierarchical relationships $pasta - isA - food$ and $dog - isA - animal$ provided by mapping S and T concepts on WordNet synsets. In this case, no $S \rightarrow T$ transformation needs to be performed. Therefore, the usage of external knowledge allows more conceptually accurate transitions. Moreover, WordNet provides concept distances necessary for edit operations, precisely reflecting semantic relationships between concepts. Then, CSED can be applied to provide the total cost of the $S \rightarrow T$ transformations.

²https://huggingface.co/models?pipeline_tag=text-to-image&sort=downloads

5. Experiments

5.1. Story Visualization

Since all semantics and **D, I** edit operations have an equal cost, we assign $d=1$ for all semantics, as well as for **D, I**. For example, deleting a color yields an edit cost of 1. Alternatively, by substituting a color with another one induces an edit cost of 2, equal to deleting the source color and then inserting the target color. The same logic applies to shape, size and material of objects.

Metric results over the best variants of selected SV models [43, 45, 46, 47] are presented in Table 1. Existing metrics (FID, Clean-FID, LPIPS, SSIM) are provided for comparison.

In general, we observe an agreement between pixel-level and conceptual metrics. This is somehow expected, since the concept extraction stage depends on pixel-level image quality, with better generated objects or semantics being more easily identifiable. Nevertheless, conceptual evaluation offers more *explainable insights*: percentages of losses per concept (Material, Size, Shape, Color) are provided, highlighting strengths and shortcomings of investigated models over different semantics. For example, higher Shape loss for all models ($> 50\%$), indicates that they synthesize objects of ambiguous shapes in most cases. On the other hand, relatively lower Size losses reveal the models’ capability to generate objects having the right size.

Table 1

Average evaluation metrics (existing and proposed, separated by vertical line) for all $L=4$ stories per M .

M	FID ↓	Clean -FID↓	LPIPS ↓	SSIM ↑	GCL ↓	GSL ↓	Material ↓	Size ↓	Shape ↓	Color ↓
[47]	41.54 ± 8.55	115.46	0.21 ± 0.05	0.71	4.97	7.01	20.83%	14.55%	56.62%	33.10%
[43]	41.45 ± 6.25	123.40	0.25 ± 0.03	0.65	11.44	15.33	30.89%	21.12%	62.34%	37.44%
[46]	41.96 ± 9.66	124.97	0.25 ± 0.08	0.67	10.95	8.06	21.45%	16.02%	56.78%	35.10%
[45]	41.80 ± 8.81	122.62	0.25 ± 0.05	0.68	8.32	11.51	25.34%	16.71%	56.83%	35.14%

We further investigate our findings by focusing on the best performing SV model of Tsakas et al. [47] according to the conceptual metrics reported in Table 1. Specifically, in Table 2 we present results of per frame **GSL**, **GCL** and losses per concept (Material, Size, Shape, Color).

Table 2

Average conceptual evaluation metrics per frame for [47].

Frame	GSC↓	GSL↓	Material↓	Size↓	Shape↓	Color↓
1st	0.00	2.25	40.00%	6.20%	58.75%	7.50%
2nd	4.35	5.66	20.00%	11.88%	57.5%	32.50%
3rd	7.12	8.25	13.33%	16.67%	57.08%	43.33%
4th	8.42	11.49	10.00%	23.44%	53.13%	49.06%

Local explanations The transparency of the proposed **SL/CL** metrics is verified by obtaining *local explanations* for [47]. Specifically, we examine edit paths for the sequences of Figure 2: the

4 leftmost images (Figure 2a) correspond to the ground truth sequence, while the 4 rightmost images (Figure 2b) denote the generated frames. Consequently, S contains concepts of 2b and T contains concepts of 2a. As presented in Table 3 (details in Appendix), a standard **R** operation for all frames is observed, suggesting transforming the material of the small brown sphere from 'rubber' to 'metallic' in order to match the ground truth. In the last frame, one more **R** operation is added, suggesting also transforming the shape of the new object from 'sphere' to 'cylinder'. The cost for each **R** operation equals to 2, equivalent for one step to remove the wrong semantic and one more step to add the right semantic. However, this cost weight can be tuned appropriately, if needed. SL for this story equals to 10, as a summary of all operation costs per frame. By observing for CL, we realize that the correct number of objects is added in every consequent frame, so that $CL_{k>1} = |\mathcal{C}| \cdot (k - 1)$, $|\mathcal{C}| = 4$ is maintained: starting from $CL_1 = p_{k=1} = 0$ for the $k=1$ frame, we verify that only one object is added, respecting that frame number should be equal to the number of objects present in it. $CL_2=4$ is expected since the object added in the $k=2$ frame contains 4 semantics. Any number lower or greater than that would indicate an abnormal behavior: $CL_{k>1} < |\mathcal{C}| \cdot (k - 1)$ marks one (or more) missing objects, while $CL_{k>1} > |\mathcal{C}| \cdot (k - 1)$ indicates one (or more) extra object generated. The desired pattern repeats for the 3rd and 4th frames. Through this analysis, the shortcomings of [47] concerning this specific image are revealed, producing a *local* explanation: The semantic *Material* needs to be examined more, as in all story frames of this example the small brown sphere is generated with the attribute 'rubber' instead of 'metallic'. In order to obtain insights regarding the model's synthesis capabilities of discrete semantics, global metrics and explanations need to be derived.




Figure 2: Ground truth vs generated CLEVR-SV story frames using [47] for $L=4$.

Table 3

Interpretable local edits for Figure 2.

Frame	Min edit path	Operation	Edit cost	Semantic	CL
1st	'rubber' → 'metallic'	R	2	Material	0
2nd	'rubber' → 'metallic'	R	2	Material	4
3rd	'rubber' → 'metallic'	R	2	Material	8
4th	{'rubber', 'sphere'} → {'metallic', 'cylinder'}	R, R	4	Material, Shape	12

Global explanations In order to assess our model's shortcomings in total, we measure **GSL** for all test images of CLEVR-SV. Therefore, we can obtain a measure of the model's inability to capture certain -discrete- semantics, either per frame or in total (Table 1). We observe that in later frames, *Material loss* decreases, even though we would expect that the problem gets

Table 4

Interpretable global edits on test set images of CLEVR-SV generated using [47].

Rules (edits)	Semantic	Support %	Antec. support%	Conseq. support%
'metallic' → 'rubber'	Material	26.77	26.77	26.77
'rubber' → 'metallic'	Material	22.05	22.05	22.83
'cylinder' → 'cube'	Shape	18.11	33.07	31.50
'cylinder' → 'sphere'	Shape	14.96	33.07	18.90

harder and harder as more objects are added, resulting in higher losses. This expected pattern is followed by *Size loss* and *Color loss*, while no certain pattern can be extracted from *Shape loss*. The high *Shape loss* imposes the need for attention mechanisms within the used GANs [28], so that long-range relationships can be captured. We can also attribute the rapid rise of *Size* and *Color* losses to consistency deficiencies within the story sequence.

GSL can also reveal patterns in the form of rules for the whole test set. We leverage the apriori algorithm [74] to extract frequent semantic combinations and rules. The 4 most common semantic edits are provided in Table 4, together with each rule’s frequency (support). The concept category (as occurring from equation 2), antecedent support (source semantic frequency), and consequent support (target semantic frequency) are also provided.

We observe that *Material* is the most common concept misconception, with both 'rubber' and 'metallic' semantics being frequently confused. *Shape* is the second most prominent misconception, with 'cylinder' appearing in the generated frames more often compared to the 'cylinder' occurrence in the conditioning; 'cube' and 'sphere' shapes are sacrificed for 'cylinder' to be generated. Since the rule support is not significantly high, with 26.77% being the maximum value, we can safely assume that the SV model of Tsakas et al. [47] is not heavily biased towards certain semantics. Nevertheless, we spot some tendency to generate the wrong material and shape, an observation that can be valuable for architectural improvements of the model.

5.2. Scene Generation

We select the first 10K samples from COCO to reduce the inference time needed to extract visual concepts using YOLO-v8 and YOLOs object detectors. COCO provides 5 descriptive sentences per sample, which are paraphrases of each other. For this reason, we only regard the 1st out of the 5 sentences as the conditioning c . We follow two separate processes for SG: actual *generation* conditioned on c and *retrieval* of caption-image pairs based on captions similar to c .

Conditional generation on COCO captions For the generation experiment, we employ pre-trained diffusion models without any further tuning, as mentioned in 4.1, which are all tested on the same conditionings c . Each of the four diffusion models required about 15 hours to synthesize 10K images using 2 T4 GPUs, therefore around 60 hours in total.

Retrieval of COCO-related captions In order to obtain considerably more images conditioned on COCO-related queries without having to spend the time and resources to run many more thousand iterations of the diffusion model, we utilized a Stable Diffusion search engine

(Lexica.art)³. The exact process we used was the following: we use c of the first 10K COCO samples as the 'query' caption. The search engine returned, for each of the 10k captions, 10 images that have been already generated by online communities with the closest input queries to our captions. This technique supplied us with 100.000 more Stable Diffusion images, accompanied by their input queries. We then compare results between web-retrieved and generated images.

Object detection We select a default threshold of $T_d = 0.6$ for detection; objects detected with confidence ≥ 0.6 are added in the generated concept set S . This threshold is experimentally defined to maintain a valid trade-off between false positive and false negative objects; in fact, since no ground truth exists, even defining false predictions is untractable without human inspection. However, our approach can provide relevant hints regarding the probability of false detection, as a higher number of **D** operations may infer higher false positive rates (irrelevant objects being detected, if T_d is too low), while more **I** operations can be correlated with higher false negative rates (relevant objects not being detected, if T_d is too strict).

Metric results For comparative reasons we present results for $T_d = 0.5, 0.6, 0.7$ in Tables 5 (YOLO-v8) & 6 (YOLOS) for generated images, and in Table 7 for web images, reporting object extraction from both object detectors. Instances colored in **blue** denote the lowest scores, which are more desirable, while the highest scores are highlighted with **red**. We present number of edits (**# I**, **# D**, **# R**), as well as the total cost for each **I**, **D**, **R** operation for all images. Mean CSED is reported as an overall metric regardless of which operation was performed more often.

Table 5

Metric results using YOLO-v8 for object detection on generated images from COCO queries.

T_d	M	# I	Cost I	# D	Cost D	# R	Cost R	Mean CSED
0.5	stable diffusion	37651	16762	1196	5655	126004	14323	35.75
	stable diffusion 2	36878	16067	1243	6301	129315	14839	36.32
	protogen base	37072	16208	1233	5944	129290	14744	35.95
	protogen 5.8	38581	17715	1195	4702	117708	13411	34.66
0.6	stable diffusion	39070	18386	1157	4042	110260	12964	34.22
	stable diffusion 2	38678	17782	1200	4514	112499	13397	34.55
	protogen base	38548	17794	1184	4270	114762	13427	34.35
	protogen 5.8	39766	19210	1134	3419	103579	12135	33.38
0.7	stable diffusion	40814	20391	1086	2681	93390	11337	32.96
	stable diffusion 2	40677	19806	1107	2938	95477	11756	33.08
	protogen base	40397	19801	1101	2820	97314	11787	32.94
	protogen 5.8	41295	20944	1039	2308	89850	10726	32.39

Regarding the selected threshold T_d , our initial hypothesis is proven to be correct: more **I** operations are realized for higher threshold $T_d=0.7$, suggesting that objects from the conditioning were not detected, while fewer **I** were performed for $T_d=0.5$. Similarly, there are more **D** operations for the lowest $T_d=0.5$, as spurious objects can be detected more easily. Additionally,

³<https://lexica.art/>

Table 6

Metric results using YOLOS for object detection on generated images from COCO queries.

T_d	M	# I	Cost I	# D	Cost D	# R	Cost R	Mean CSED
0.5	stable diffusion	26302	9032	1382	44189	197623	21097	68.25
	stable diffusion 2	26684	8832	1403	43459	192198	21082	68.05
	protogen base	26887	8966	1404	44406	193327	21035	68.81
	protogen 5.8	28880	10367	1373	34996	189677	19858	60.45
0.6	stable diffusion	27963	9920	1373	33891	188395	20286	60.10
	stable diffusion 2	28145	9662	1394	33933	182767	20322	60.36
	protogen base	28499	9845	1394	34167	185217	20224	60.63
	protogen 5.8	30545	11330	1364	27218	179947	18963	54.13
0.7	stable diffusion	29998	10985	1357	24956	177213	19319	52.51
	stable diffusion 2	29831	10657	1347	25492	172860	19409	53.14
	protogen base	29866	10790	1346	25255	175495	19350	52.98
	protogen 5.8	28880	10367	1373	34996	189677	19858	60.45

Table 7

Metric results for web-retrieved Stable Diffusion images on similar queries to COCO.

T_d	Obj. detector	# I	Cost I	# D	Cost D	# R	Cost R	Mean CSED
0.5	YOLO-v8	186775	857448	1343	52247	1353479	224350	75.87
	YOLOS	163628	605321	1487	421525	2469635	473331	106.41
0.6	YOLO-v8	190047	891454	1317	37418	1174012	190928	73.74
	YOLOS	167576	646112	1467	308346	2303966	432851	98.06
0.7	YOLO-v8	193663	929183	1236	25388	982259	154063	71.81
	YOLOS	171778	688942	1449	214928	2115779	390304	90.56

more **R** operations are needed for lower thresholds, which is also expected, since more objects are extracted and added to the S set. As for object detectors, results using YOLO-v8 are very homogeneous, indicating that the models under investigation follow a rather predictable behavior irrespectively of T_d . Protogen 5.8 consistently yields the lowest mean CSED score, denoting cheaper transitions for all thresholds. This observation slightly changes for $T_d=0.7$ and YOLOS object detector (Table 6), for which, surprisingly, protogen 5.8 produces the more expensive transitions. By comparing Tables 5 & 6, YOLOS results in higher mean CSED, less **I** operations, significantly more expensive **D** operations (even though the number of **D** operations is not substantially larger), as well as more and expensive **R** operations. Therefore, we can safely assume that YOLOS is comparatively more sensitive in detecting more objects, which may induce some noise in the detection process. All these results will become more interpretable should we delve into the explanations accompanying the evaluation. The patterns arising from evaluating generated images are also supported in Table 7 findings, verifying the threshold hypothesis, as well as the increased sensitivity of YOLOS. Nevertheless, web-retrieved images seem to miss objects mentioned in the query, as proven by the large number of **I** and **R** operations.

Local explanations provide edit paths based on the **I**, **D**, **R** operations realized for a specific generated image. For this reason, we employ a scene depicted in Figure 3.




Figure 3: An image sample generated by Stable Diffusion 2 to extract local explanations.

According to YOLO-v8 with the default threshold $T_d=0.6$, the generated concepts are $S=\{\text{'car'}$, 'car' , 'traffic light' , 'car' , $\text{'stop sign'}\}$, and ground truth concepts are $T=\{\text{'light'}$, $\text{'buildings'}\}$. The edit operations of total minimum cost 59.00 for this $S \rightarrow T$ transformation are:

I: $\{\}$

D: $\{\text{'car'}$, 'car' , $\text{'car'}\}$

R: $\{\text{'traffic light'} \rightarrow \text{'light'}$, $\text{'stop sign'} \rightarrow \text{'buildings'}\}$

When using YOLOS, the generated concepts are $S=\{\text{'car'}$, 'traffic light' , 'car' , 'stop sign' , 'traffic light' , 'car' , 'traffic light' , 'traffic light' , 'traffic light' , 'traffic light' , 'traffic light' , 'traffic light' , 'traffic light' , 'car' , 'traffic light' , 'traffic light' , 'traffic light' , 'traffic light' , 'traffic light' , 'car' , 'car' , 'traffic light' , $\text{'traffic light'}\}$, and the ground truth ones are $T=\{\text{'light'}$, $\text{'buildings'}\}$. By visually inspecting the image, YOLOS clearly overestimates the actual objects present, inducing noise in the generated concept set S . Nevertheless, our evaluation strategy successfully captures this overestimation, by suggesting the deletion of multiple concepts. Specifically, we obtain the following transformations of total cost 104.04:

I: $\{\}$

D: $\{\text{'car'}$, 'traffic light' , 'car' , 'traffic light' , 'car' , 'traffic light' , 'traffic light' , 'traffic light' , 'traffic light' , 'traffic light' , 'traffic light' , 'traffic light' , 'traffic light' , 'traffic light' , 'traffic light' , 'traffic light' , 'car' , 'car' , 'traffic light' , $\text{'traffic light'}\}$

R: $\{\text{'stop sign'} \rightarrow \text{'light'}$, $\text{'traffic light'} \rightarrow \text{'buildings'}\}$

Global explanations for all images are presented in Table 8 for **I**, **D** edits and Table 9 for **R** edits. Results only involve YOLO-v8 extracted concepts, as YOLOS results in an overwhelming number of detected instances. Top-3 results are demonstrated, i.e. the 3 most frequent insertion, deletions and replacements. **I** and **D** refers to concepts inserted or deleted respectively, while Freq **I**, **D** denotes how many times a specific concepts was inserted or deleted within all images. **I**, **D** support indicates the frequency a specific edit happens among all **I**, **D** edits respectively. As for **R**, support denotes the frequency of a transformation rule among all produced rules.

Table 8
Global explanations (**I** and **D** edits) for YOLO-v8 extracted concepts.

T_d	M	I	Freq I	I support	D	Freq D	D support
0.5	stable diffusion	street	264	1.57%	person	2075	36.69%
		table	250	1.49%	sheep	363	6.42%
		tennis	247	1.47%	car	252	4.46%
	stable diffusion 2	tennis	253	1.57%	person	2177	34.55%
		street	242	1.51%	sheep	466	7.40%
		table	237	1.48%	car	313	4.97%
	protogen base	tennis	247	1.52%	person	2281	38.37%
		street	244	1.51%	sheep	317	5.33%
		table	229	1.41%	car	311	5.23%
	protogen 5.8	table	270	1.52%	person	1564	33.26%
		tennis	265	1.50%	car	261	5.55%
		street	241	1.36%	umbrella	251	5.34%
0.6	stable diffusion	street	290	1.58%	person	1572	38.89%
		table	281	1.53%	sheep	311	7.69%
		tennis	259	1.41%	car	158	3.91%
	stable diffusion 2	table	274	1.54%	person	1656	36.69%
		street	269	1.51%	sheep	376	8.33%
		tennis	264	1.48%	car	203	4.50%
	protogen base	street	268	1.51%	person	1717	40.21%
		table	261	1.47%	sheep	254	5.95%
		tennis	255	1.43%	car	197	4.61%
	protogen 5.8	table	303	1.58%	person	1220	35.68%
		tennis	278	1.45%	sheep	198	5.79%
		street	274	1.43%	umbrella	176	5.15%
0.7	stable diffusion	table	322	1.58%	person	1075	40.10%
		street	316	1.55%	sheep	254	9.47%
		tennis	268	1.31%	donut	122	4.55%
	stable diffusion 2	table	313	1.58%	person	1134	38.60%
		street	301	1.52%	sheep	291	9.90%
		tennis	267	1.35%	donut	111	3.78%
	protogen base	street	300	1.52%	person	1189	42.16%
		table	289	1.46%	sheep	188	6.67%
		tennis	262	1.32%	umbrella	143	5.07%
	protogen 5.8	table	330	1.58%	person	884	38.30%
		street	299	1.43%	sheep	152	6.59%
		tennis	287	1.37%	umbrella	130	5.63%

We can observe an obvious agreement between models; **I** edits include 'street', 'tennis' and 'table' concepts. It seems that the selected M cannot efficiently generate the **I** concepts, or generated concepts are of low visual quality, so that their detection is not feasible with $T_d=0.5, 0.6, 0.7$. **D** edits mainly contain 'person', 'sheep', 'car', 'umbrella', 'donut' concepts, indicating some bias towards generating spurious instances of those concept categories. Finally, **R** edits

refer to transforming 'person' to 'people', 'man' or 'woman'. Since 'person' is a YOLO category incorporating both genders, such transformations are somehow expected.

Table 9
Global explanations (**R** edits) for YOLO-v8 extracted concepts.

T_d	M	R	Freq R	R sup- port	M	R	Freq R	R sup- port
0.5	stable diffusion	person → man	1090	7.61%	stable diffusion	person → man	1115	7.51%
		person → people	520	3.63%		person → people	551	3.71%
		person → woman	499	3.48%		2	person → woman	511
	protogen base	person → man	1101	7.47%	protogen	person → man	1061	7.91%
		person → people	507	3.44%		person → woman	476	3.55%
		person → woman	500	3.39%		5.8	person → people	441
0.6	stable diffusion	person → man	1065	8.22%	stable diffusion	person → man	1087	8.11%
		person → people	503	3.88%		person → people	536	4.00%
		person → woman	481	3.71%		2	person → woman	482
	protogen base	person → man	1080	8.04%	protogen	person → man	1035	8.53%
		person → people	494	3.68%		person → woman	449	3.70%
		person → woman	485	3.61%		5.8	person → people	431
0.7	stable diffusion	person → man	1022	9.01%	stable diffusion	person → man	1033	8.79%
		person → people	473	4.17%		person → people	508	4.32%
		person → woman	458	4.04%		2	person → woman	441
	protogen base	person → man	1054	8.94%	protogen	person → man	989	9.22%
		person → woman	461	3.91%		person → woman	419	3.91%
		person → people	446	3.78%		5.8	person → people	408

6. Conclusion

Conceptual approaches in generative evaluation is an underexplored field, which can provide some novel insights regarding model quality and explainability of results. In our work, we propose a knowledge-driven explainable evaluation framework that suggests which concepts should be added, removed, or replaced for a generated image to approach its conditioning. Results on competitive tasks such as Story Visualization and Scene Generation illustrate the merits of such an approach, highlighting concepts that models cannot generate, or model biases towards generating excessive numbers of specific concept categories. As future work, we plan to expand our approach to other models and tasks and also incorporate alternative knowledge sources to examine how the produced edit paths conceptually deviate from the current ones.

Acknowledgments

The research work was supported by the Hellenic Foundation for Research and Innovation (HFRI) under the 3rd Call for HFRI PhD Fellowships (Fellowship Number 5537).

References

- [1] J. Ho, A. Jain, P. Abbeel, Denoising diffusion probabilistic models, 2020. URL: <https://arxiv.org/abs/2006.11239>. doi:10.48550/ARXIV.2006.11239.
- [2] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, High-resolution image synthesis with latent diffusion models, in: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), 2022, pp. 10684–10695.
- [3] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, I. Sutskever, Zero-shot text-to-image generation, 2021. arXiv:2102.12092.
- [4] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, M. Chen, Hierarchical text-conditional image generation with clip latents, 2022. URL: <https://arxiv.org/abs/2204.06125>. doi:10.48550/ARXIV.2204.06125.
- [5] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. S. Ghasemipour, B. K. Ayan, S. S. Mahdavi, R. G. Lopes, T. Salimans, J. Ho, D. J. Fleet, M. Norouzi, Photorealistic text-to-image diffusion models with deep language understanding, 2022. arXiv:2205.11487.
- [6] A. Blattmann, R. Rombach, K. Oktay, J. Müller, B. Ommer, Semi-parametric neural image synthesis, 2022. URL: <https://arxiv.org/abs/2204.11824>. doi:10.48550/ARXIV.2204.11824.
- [7] J. Kim, M. Lee, Class-continuous conditional generative neural radiance field, 2023. URL: <https://arxiv.org/abs/2301.00950>. doi:10.48550/ARXIV.2301.00950.
- [8] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, Y. Bengio, Generative adversarial nets, in: Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, K. Q. Weinberger (Eds.), Advances in Neural Information Processing Systems, volume 27, Curran Associates, Inc., 2014. URL: <https://proceedings.neurips.cc/paper/2014/file/5ca3e9b122f61f8f06494c97b1afccf3-Paper.pdf>.
- [9] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, X. Chen, X. Chen, Improved techniques for training gans, in: D. Lee, M. Sugiyama, U. Luxburg, I. Guyon, R. Garnett (Eds.), Advances in Neural Information Processing Systems, volume 29, Curran Associates, Inc., 2016. URL: <https://proceedings.neurips.cc/paper/2016/file/8a3363abe792db2d8761d6403605aeb7-Paper.pdf>.
- [10] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, S. Hochreiter, Gans trained by a two time-scale update rule converge to a local nash equilibrium, in: I. Guyon, U. V. Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, R. Garnett (Eds.), Advances in Neural Information Processing Systems, volume 30, Curran Associates, Inc., 2017. URL: <https://proceedings.neurips.cc/paper/2017/file/8a1d694707eb0fefe65871369074926d-Paper.pdf>.
- [11] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, O. Wang, The unreasonable effectiveness of deep features as a perceptual metric, in: CVPR, 2018.
- [12] G. Parmar, R. Zhang, J.-Y. Zhu, On aliased resizing and surprising subtleties in gan evaluation, in: CVPR, 2022.
- [13] A. Borji, Pros and cons of gan evaluation measures: New developments, Computer Vision and Image Understanding 215 (2022) 103329. URL: <https://www.sciencedirect.com/science/article/pii/S1077314221001685>. doi:<https://doi.org/10.1016/j.cviu.2021.103329>.
- [14] M. Soloveitchik, T. Diskin, E. Morin, A. Wiesel, Conditional frechet inception distance,

2021. URL: <https://arxiv.org/abs/2103.11521>. doi:10.48550/ARXIV.2103.11521.
- [15] Y. Benny, T. Galanti, S. Benaim, L. Wolf, Evaluation metrics for conditional image generation, *International Journal of Computer Vision* 129 (2021) 1712–1731. URL: <https://doi.org/10.1007/s11263-020-01424-w>. doi:10.1007/s11263-020-01424-w.
- [16] K. Abhishek, D. Kamath, Attribution-based xai methods in computer vision: A review, 2022. URL: <https://arxiv.org/abs/2211.14736>. doi:10.48550/ARXIV.2211.14736.
- [17] V. Buhrmester, D. Münch, M. Arens, Analysis of explainers of black box deep neural networks for computer vision: A survey, *Machine Learning and Knowledge Extraction* 3 (2021) 966–989. URL: <https://www.mdpi.com/2504-4990/3/4/48>. doi:10.3390/make3040048.
- [18] V. Nagisetty, L. Graves, J. Scott, V. Ganesh, xai-gan: Enhancing generative adversarial networks via explainable ai systems, 2020. URL: <https://arxiv.org/abs/2002.10438>. doi:10.48550/ARXIV.2002.10438.
- [19] A. Genovese, V. Piuri, F. Scotti, Towards explainable face aging with generative adversarial networks, in: 2019 IEEE International Conference on Image Processing (ICIP), 2019, pp. 3806–3810. doi:10.1109/ICIP.2019.8803616.
- [20] J. Kim, H. Park, Limited discriminator gan using explainable ai model for overfitting problem, *ICT Express* (2022). URL: <https://www.sciencedirect.com/science/article/pii/S240595952100179X>. doi:<https://doi.org/10.1016/j.ictex.2021.12.014>.
- [21] C. Leiter, P. Lertvittayakumjorn, M. Fomicheva, W. Zhao, Y. Gao, S. Eger, Towards explainable evaluation metrics for natural language generation, 2022. URL: <https://arxiv.org/abs/2203.11131>. doi:10.48550/ARXIV.2203.11131.
- [22] J. Opitz, A. Frank, Towards a decomposable metric for explainable evaluation of text generation from AMR, in: Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, Association for Computational Linguistics, Online, 2021, pp. 1504–1518. URL: <https://aclanthology.org/2021.eacl-main.129>. doi:10.18653/v1/2021.eacl-main.129.
- [23] M. Lymperaio, G. Manoliadis, O. Menis Mastromichalakis, E. G. Dervakos, G. Stamou, Towards explainable evaluation of language models on the semantic similarity of visual concepts, in: Proceedings of the 29th International Conference on Computational Linguistics, International Committee on Computational Linguistics, Gyeongju, Republic of Korea, 2022, pp. 3639–3658. URL: <https://aclanthology.org/2022.coling-1.321>.
- [24] G. Filandrianos, K. Thomas, E. Dervakos, G. Stamou, Conceptual edits as counterfactual explanations, in: Proceedings of the AAAI 2022 Spring Symposium on Machine Learning and Knowledge Engineering for Hybrid Intelligence (AAAI-MAKE 2022), Stanford University, Palo Alto, California, USA, 2022.
- [25] M. Mirza, S. Osindero, Conditional generative adversarial nets, *CoRR* abs/1411.1784 (2014). URL: <http://arxiv.org/abs/1411.1784>. arXiv:1411.1784.
- [26] A. Odena, C. Olah, J. Shlens, Conditional image synthesis with auxiliary classifier gans, 2017. arXiv:1610.09585.
- [27] T. Miyato, M. Koyama, cgans with projection discriminator, 2018. arXiv:1802.05637.
- [28] H. Zhang, I. Goodfellow, D. Metaxas, A. Odena, Self-attention generative adversarial networks, 2019. arXiv:1805.08318.
- [29] T. Miyato, T. Kataoka, M. Koyama, Y. Yoshida, Spectral normalization for generative adversarial networks, 2018. arXiv:1802.05957.

- [30] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, S. Hochreiter, Gans trained by a two time-scale update rule converge to a local nash equilibrium, 2018. [arXiv:1706.08500](https://arxiv.org/abs/1706.08500).
- [31] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, B. Ommer, High-resolution image synthesis with latent diffusion models, 2021. [arXiv:2112.10752](https://arxiv.org/abs/2112.10752).
- [32] A. Radford, J. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, I. Sutskever, Learning transferable visual models from natural language supervision, 2021.
- [33] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. S. Ghasemipour, B. K. Ayan, S. S. Mahdavi, R. G. Lopes, T. Salimans, J. Ho, D. J. Fleet, M. Norouzi, Photorealistic text-to-image diffusion models with deep language understanding, 2022. URL: <https://arxiv.org/abs/2205.11487>. doi:10.48550/ARXIV.2205.11487.
- [34] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, P. J. Liu, Exploring the limits of transfer learning with a unified text-to-text transformer, *Journal of Machine Learning Research* 21 (2020) 1–67. URL: <http://jmlr.org/papers/v21/20-074.html>.
- [35] N. Ruiz, Y. Li, V. Jampani, Y. Pritch, M. Rubinstein, K. Aberman, Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation, 2022. URL: <https://arxiv.org/abs/2208.12242>. doi:10.48550/ARXIV.2208.12242.
- [36] S. E. Reed, Z. Akata, X. Yan, L. Logeswaran, B. Schiele, H. Lee, Generative adversarial text to image synthesis, *CoRR abs/1605.05396* (2016). URL: <http://arxiv.org/abs/1605.05396>. [arXiv:1605.05396](https://arxiv.org/abs/1605.05396).
- [37] S. E. Reed, Z. Akata, S. Mohan, S. Tenka, B. Schiele, H. Lee, Learning what and where to draw, *CoRR abs/1610.02454* (2016). URL: <http://arxiv.org/abs/1610.02454>. [arXiv:1610.02454](https://arxiv.org/abs/1610.02454).
- [38] H. Zhang, T. Xu, H. Li, Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks, 2017, pp. 5908–5916. doi:10.1109/ICCV.2017.629.
- [39] H. Zhang, T. Xu, H. Li, S. Zhang, X. Wang, X. Huang, D. Metaxas, Stackgan++: Realistic image synthesis with stacked generative adversarial networks, 2018. [arXiv:1710.10916](https://arxiv.org/abs/1710.10916).
- [40] T. Xu, P. Zhang, Q. Huang, H. Zhang, Z. Gan, X. Huang, X. He, Attngan: Fine-grained text to image generation with attentional generative adversarial networks, in: *CVPR 2018*, 2018.
- [41] H. Tan, X. Liu, X. Li, Y. Zhang, B. Yin, Semantics-enhanced adversarial nets for text-to-image synthesis, in: *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019, pp. 10500–10509. doi:10.1109/ICCV.2019.01060.
- [42] M. Zhu, P. Pan, W. Chen, Y. Yang, DM-GAN: dynamic memory generative adversarial networks for text-to-image synthesis, *CoRR abs/1904.01310* (2019). URL: <http://arxiv.org/abs/1904.01310>. [arXiv:1904.01310](https://arxiv.org/abs/1904.01310).
- [43] Y. Li, Z. Gan, Y. Shen, J. Liu, Y. Cheng, Y. Wu, L. Carin, D. E. Carlson, J. Gao, Storygan: A sequential conditional gan for story visualization, *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2019) 6322–6331.
- [44] C. Li, L. Kong, Z. Zhou, Improved-storygan for sequential images visualization, *Journal of Visual Communication and Image Representation* 73 (2020) 102956. URL: <https://www.sciencedirect.com/science/article/pii/S1047320320301826>. doi:<https://doi.org/10.1016/j.jvcir.2020.102956>.
- [45] A. Maharana, D. Hannan, M. Bansal, Improving generation and evaluation of visual stories

- via semantic consistency, ArXiv abs/2105.10026 (2021).
- [46] A. Maharana, M. Bansal, Integrating visuospatial, linguistic, and commonsense structure into story visualization, ArXiv abs/2110.10834 (2021).
 - [47] N. Tsakas, M. Lymperaïou, G. Filandrianos, G. Stamou, An impartial transformer for story visualization, 2023. URL: <https://arxiv.org/abs/2301.03563>. doi:10.48550/ARXIV.2301.03563.
 - [48] A. Horé, D. Ziou, Image quality metrics: Psnr vs. ssim, in: 2010 20th International Conference on Pattern Recognition, 2010, pp. 2366–2369. doi:10.1109/ICPR.2010.579.
 - [49] D. Bau, J.-Y. Zhu, J. Wulff, W. Peebles, H. Strobel, B. Zhou, A. Torralba, Seeing what a gan cannot generate, 2019. URL: <https://arxiv.org/abs/1910.11626>. doi:10.48550/ARXIV.1910.11626.
 - [50] Y. Shen, J. Gu, X. Tang, B. Zhou, Interpreting the latent space of gans for semantic face editing, 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019) 9240–9249.
 - [51] Y. Shen, C. Yang, X. Tang, B. Zhou, Interfacegan: Interpreting the disentangled face representation learned by gans, IEEE Transactions on Pattern Analysis and Machine Intelligence 44 (2020) 2004–2018.
 - [52] L. Chai, J. Wulff, P. Isola, Using latent space regression to analyze and leverage compositionality in gans, ArXiv abs/2103.10426 (2021).
 - [53] J. Wen, F. Benitez-Quiroz, Q. Feng, A. M. Martinez, Diamond in the rough: Improving image realism by traversing the gan latent space, ArXiv abs/2104.05518 (2021).
 - [54] M. Nauta, J. Trienes, S. Pathak, E. Nguyen, M. Peters, Y. Schmitt, J. Schlötterer, M. van Keulen, C. Seifert, From anecdotal evidence to quantitative evaluation methods: A systematic review on evaluating explainable ai, 2022. URL: <https://arxiv.org/abs/2201.08164>. doi:10.48550/ARXIV.2201.08164.
 - [55] W. Zhao, S. Oyama, M. Kurihara, Generating natural counterfactual visual explanations, in: C. Bessiere (Ed.), Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20, International Joint Conferences on Artificial Intelligence Organization, 2020, pp. 5204–5205. URL: <https://doi.org/10.24963/ijcai.2020/742>. doi:10.24963/ijcai.2020/742, doctoral Consortium.
 - [56] R. Poyiadzi, K. Sokol, R. Santos-Rodríguez, T. D. Bie, P. A. Flach, Face: Feasible and actionable counterfactual explanations, Proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (2020).
 - [57] Y. Goyal, Z. Wu, J. Ernst, D. Batra, D. Parikh, S. Lee, Counterfactual visual explanations, in: K. Chaudhuri, R. Salakhutdinov (Eds.), Proceedings of the 36th International Conference on Machine Learning, volume 97 of *Proceedings of Machine Learning Research*, PMLR, 2019, pp. 2376–2384. URL: <https://proceedings.mlr.press/v97/goyal19a.html>.
 - [58] C. Fellbaum, Wordnet: An electronic lexical database (1998).
 - [59] K. Browne, B. Swift, Semantics and explanation: why counterfactual explanations produce adversarial examples in deep neural networks, ArXiv abs/2012.10076 (2020).
 - [60] J. Johnson, B. Hariharan, L. van der Maaten, L. Fei-Fei, C. L. Zitnick, R. B. Girshick, Clevr: A diagnostic dataset for compositional language and elementary visual reasoning, 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2017) 1988–1997.
 - [61] T.-Y. Lin, M. Maire, S. Belongie, L. Bourdev, R. Girshick, J. Hays, P. Perona, D. Ramanan, C. L.

- Zitnick, P. Dollár, Microsoft coco: Common objects in context, 2015. [arXiv:1405.0312](https://arxiv.org/abs/1405.0312).
- [62] Stable diffusion v1.4, Huggingface. URL: <https://huggingface.co/CompVis/stable-diffusion-v1-4>.
- [63] Stable diffusion 2 base, Huggingface. URL: <https://huggingface.co/stabilityai/stable-diffusion-2-base>.
- [64] Protogen x3.4, Huggingface. URL: https://huggingface.co/darkstorm2150/Protogen_x3.4_Official_Release.
- [65] Protogen x5.8, Huggingface. URL: https://huggingface.co/darkstorm2150/Protogen_x5.8_Official_Release.
- [66] T. Park, M.-Y. Liu, T.-C. Wang, J.-Y. Zhu, Semantic image synthesis with spatially-adaptive normalization, 2019. [arXiv:1903.07291](https://arxiv.org/abs/1903.07291).
- [67] A. Zareian, Z. Wang, H. You, S.-F. Chang, Learning visual commonsense for robust scene graph generation, 2020. [arXiv:2006.09623](https://arxiv.org/abs/2006.09623).
- [68] K. Ma, B. Zhao, L. Sigal, Attribute-guided image generation from layout, 2020. [arXiv:2008.11932](https://arxiv.org/abs/2008.11932).
- [69] Y. Li, Y. Cheng, Z. Gan, L. Yu, L. Wang, J. Liu, Bachgan: High-resolution image synthesis from salient object layout, 2020. [arXiv:2003.11690](https://arxiv.org/abs/2003.11690).
- [70] W. Sun, T. Wu, Learning layout and style reconfigurable gans for controllable image synthesis, 2021. [arXiv:2003.11571](https://arxiv.org/abs/2003.11571).
- [71] G. Jocher, A. Chaurasia, J. Qiu, Yolo by ultralytics, 2023. URL: <https://github.com/ultralytics/ultralytics>.
- [72] Y. Fang, B. Liao, X. Wang, J. Fang, J. Qi, R. Wu, J. Niu, W. Liu, You only look at one sequence: Rethinking transformer in vision through object detection, CoRR abs/2106.00666 (2021). URL: <https://arxiv.org/abs/2106.00666>. [arXiv:2106.00666](https://arxiv.org/abs/2106.00666).
- [73] Industrial-strength natural language processing, spaCy. URL: <https://spacy.io/>.
- [74] R. Agrawal, R. Srikant, Fast algorithms for mining association rules, in: Proc. of 20th Intl. Conf. on VLDB, 1994, pp. 487–499.

A. Online Resources

The following models were used for image generation:

- Huggingface text-to-image generation models.
- Stable diffusion v1.4
- Stable diffusion 2
- Protogen x3.4
- Protogen x5.8

The following web-page was used for retrieving already generated images with Stable-Diffusion:

- <https://lexica.art/>

The following models were used for object detection:

- Huggingface YOLOS-base
- Huggingface YOLO-v8

B. Detailed local SV example

In this section, we are going to provide a detailed analysis of the local properties of CSED per frame for SV regarding the following sequence of Figure 2. This is an example of medium difficulty, according to the analysis followed in [47], as in the 4th frame the blue cylinder overlaps with the blue cube. We are going to compare the ground truth semantics of the sequence, corresponding to the ground truth frames (Figure 2a) with the generated semantics, corresponding to the generated frames (Figure 2b). We, therefore, obtain the following results:

Frame k=1

Ground truth semantics: {[small, brown, **metallic**, sphere]}

Generated semantics: {[small, brown, **rubber**, sphere]}

The two sequences differ by the highlighted semantic in the 3rd position: while the ground truth semantic is 'metallic', the generated is 'rubber', therefore $CSED_{k=1}$ proposes the replacement operation 'rubber' \rightarrow 'metallic' with Edit cost = 2 = $CSED_{k=1}$ in order for the generated sequence to become identical to the ground truth one. Moreover, as the transformation is an instance involving the *Material* semantic, one more generation failure is added to the *Material Loss* counter, which is going to provide global explanations regarding semantic synthesis failures for all test set frames.

For the 1st frame, Consistency Loss (CL) for the generated sequence is $CL_{k=1} = 0$, since there are $|\mathcal{C}|=4$ semantics in total (*Material, Size, Shape, Color*), and 1 object containing $T = 4$ semantics is placed in k=1 position in the sequence: $CL_{k=1} = p_{k=1} = |T| - |\mathcal{C}| \cdot k = 4 - 4 = 0$.

Frame k=2

Ground truth semantics: {[small, brown, **metallic**, sphere], [small, brown, metallic, sphere]}

Generated semantics: {[small, brown, **rubber**, sphere], [small, brown, metallic, sphere]}

There is a difference in the semantic of the 3rd position, highlighted in bold: while the ground truth semantic is 'metallic', the generated is 'rubber', therefore $CSED_{k=2}$ proposes the replacement operation 'rubber' \rightarrow 'metallic' with Edit cost = 2 = $CSED_{k=2}$. Moreover, as the transformation is an instance involving the *Material* semantic, one more generation failure is added to the *Material Loss* counter.

In the same time, CL will inevitably increase just by adding one more object containing $|\mathcal{C}| = 4$ semantics. Therefore, the minimum increase of CL for CLEVR-SV when one object is added can be 4. Other than that, if there are more inconsistencies between $k = 1$ and $k = 2$ generated frames, $CL_{k=2}$ will increase. Therefore, we compare $k = 1$ generated sequence $T = S_{k=1} = \{[small, brown, rubber, sphere]\}$ with the $k = 2$ generated sequence $S = S_k = \{[small, brown, rubber, sphere], [small, brown, **metallic**, sphere]\}$, where no extra differences are spotted. By applying equation 5 for k=2 we obtain:

$$CL_{k=2} = p_{k=1} + D(S_{k=2}, T_{k=2}) = 0 + \mathbf{I}\{small, brown, metallic, sphere\} = 0 + 4 = 4$$

Frame k=3

Ground truth semantics: {[small, brown, **metallic**, sphere], [small, brown, metallic, sphere], [large, blue, rubber, cube]}

Generated semantics: {[small, brown, **rubber**, sphere], [small, brown, metallic, sphere], [large, blue, rubber, cube]} The difference in the 3rd position semantic remains, therefore $CSED_{k=3}$ proposes the replacement operation 'rubber' \rightarrow 'metallic' with Edit cost = 2 = $CSED_{k=3}$. Moreover, as the transformation is an instance involving the *Material* semantic, one more generation failure is added to the *Material Loss* counter.

CL will take into account the comparison between $T = 2$ generated sequence [small, brown, rubber, sphere, small, brown, metallic, sphere] and $k = 3$ generated sequence {[small, brown, rubber, sphere], [small, brown, metallic, sphere], [**large, blue, rubber, cube**]}, which only differ by the addition of the *large, blue, rubber, cube* in the third frame, thus yielding:

$$CL_{k=3} = p_{k=1} + D(S_{k=2}, T_{k=2}) + D(S_{k=3}, T_{k=3}) = 0 + 4 + \mathbf{I}\{large, blue, rubber, cube\} = 0+4+4 = 8$$

Frame k=4

Ground truth semantics: {[small, brown, **metallic**, sphere], [small, brown, metallic, sphere], [large, blue, rubber, cube], [large, blue, metallic, **cylinder**]}

Generated semantics: {[small, brown, **rubber**, sphere], [small, brown, metallic, sphere], [large, blue, rubber, cube], [large, blue, metallic, **sphere**]}

Apart from the difference in the 3rd position semantic, for which $CSED_{k=4}$ proposes the replacement operation 'rubber' \rightarrow 'metallic' with Edit cost = 2, there is also one difference in the last position semantic, indicating the transformation 'sphere' \rightarrow 'cylinder' with Edit cost = 2. By aggregating the two transformations together, we obtain the total transformation for $k = 4$: {'rubber', 'sphere'} \rightarrow {'metallic', 'cylinder'} with Edit cost = 4 = $CSED_{k=4}$. Counters for *Material Loss* and *Shape Loss* will increase by 1 each.

For CL, the sequences corresponding to $k = 3$ generated sequence {[small, brown, rubber, sphere], [small, brown, metallic, sphere], [large, blue, rubber, cube]} and $k = 4$ generated sequence {[small, brown, rubber, sphere], [small, brown, metallic, sphere], [large, blue, rubber, cube], [**large, blue, metallic, sphere**]}, which only differ by the addition of the *large, blue, metallic, sphere* item. Therefore $CL_{T=4} = 4$.

$CL_{k=4} = p_{k=1} + D(S_{k=2}, T_{k=2}) + D(S_{k=3}, T_{k=3}) + D(S_{k=4}, T_{k=4}) = 0 + 4 + 4 + \mathbf{I}\{large, blue, metallic, sphere\} = 0+4+4+4 = 12$ By aggregating results, *Story Loss (SL)* as a sum of per frame CSED costs will be:

$$SL = 2 + 2 + 2 + 4 = 10$$

and by averaging SL on all $L = 4$ frames according to equation 7:

$$Average\ SL = \frac{1}{k} SL = 10/4 = 2.5$$

For consistency, we follow equation 8:

$$Average\ CL = \frac{p_{k=1}}{k} + \frac{1}{k} \sum_{k=1}^{k=L} [CL_{k>1} \neq |C| \cdot (k - 1)] = 0 + 0 = 0$$

The generated story of Figure 2 is **fully consistent** as the *Average CL* equals to 0, which is the ideal case. Therefore, no semantics are inserted, deleted, or altered within the generated

sequence. It is however interesting that CL cannot capture the **faithfulness** error between the new item inserted in the 4th frame: while the ground truth item is a *large, blue, metallic, cylinder*, the generated sequence inserts a *large, blue, metallic, sphere*, but CL does not penalize more the difference in the semantic of the last position. On the contrary, SL is responsible to penalize for this error. Of course, the opposite scenario could be applicable in a different example, where CL would indicate an error that SL could not capture. This observation concludes that both metrics can be important, with SL focusing on faithfulness between ground truth and generated stories, while CL focuses on consistency between consequently generated frames. The better the model, the *lower* both metrics should be on the global level.