

Artificial Intelligence for Geometry-Based Feature Extraction, Analysis and Synthesis in Artistic Images: A Survey

Mridula Vijendran¹, Jingjing Deng¹, Shuang Chen¹,
Edmond S. L. Ho², Hubert P. H. Shum^{1*}

¹Department of Computer Science, Durham University, Durham,
DH1 3LE, United Kingdom.

²School of Computer Science, University of Glasgow, Glasgow,
G12 8RZ, United Kingdom.

*Corresponding author(s). E-mail(s): hubert.shum@durham.ac.uk;
Contributing authors: mridula.vijendran@durham.ac.uk;
jingjing.deng@durham.ac.uk; shuang.chen@durham.ac.uk;
shu-lim.ho@glasgow.ac.uk;

Abstract

Artificial Intelligence significantly enhances the visual art industry by analyzing, identifying and generating digitized artistic images. This review highlights the substantial benefits of integrating geometric data into AI models, addressing challenges such as high inter-class variations, domain gaps, and the separation of style from content by incorporating geometric information. Models not only improve AI-generated graphics synthesis quality, but also effectively distinguish between style and content, utilizing inherent model biases and shared data traits. We explore methods like geometric data extraction from artistic images, the impact on human perception, and its use in discriminative tasks. The review also discusses the potential for improving data quality through innovative annotation techniques and the use of geometric data to enhance model adaptability and output refinement. Overall, incorporating geometric guidance boosts model performance in classification and synthesis tasks, providing crucial insights for future AI applications in the visual arts domain.

Keywords: Artificial intelligence, machine learning, feature extraction, geometrical analysis, content synthesis.

1 Introduction

Artificial intelligence (AI) techniques find use in the art industry for tasks such as 3D scan analysis, art recommendation systems, identification of art design principles, deconstructivism art generation with fragment models [1, 2]. These techniques mainly involve three key processes: extraction, analysis, and synthesis. Extraction methods classify paintings based on style, identify and authenticate artwork, and provide exhibit and tour information to establishments like museums [3] and historic cathedrals [4] to enhance their visitors’ experience. Analysis techniques interpret and facilitate searching and comparing art elements such as geometric patterns of compositional elements across multiple scales in art collections [5]. Synthesis methods are used for scanning and enhancing the details of artifacts in the Cultural Heritage field [6, 7] and deal with the preservation, documentation and collection of historical and cultural objects. They classify paintings based on style, identify and authenticate artwork, and provide virtual access to historic cathedrals to enhance their consumers’ experience. It is used for scanning and enhancing the details of artifacts in the Cultural Heritage field [6–8] which deals with the preservation, documentation and collection of historical and cultural objects.

This paper explores a broad spectrum of AI techniques applied to artworks, encompassing both deep learning and non-deep learning methods. While deep learning models, such as CNNs and GANs, have shown significant success in tasks like object detection and image synthesis, we also discuss traditional methods, including Deformable Part Models (DPM), Histogram of Oriented Gradients (HoG), and Thin Plate Spline (TPS) interpolation. These non-deep learning techniques are crucial for understanding specific geometric features and enhancing overall model performance when integrated with deep learning approaches.

Learning from Art datasets using artificial intelligence models is challenging due to the generally smaller dataset size and larger inter-class variations [9, 10], as well as incomplete or inaccurate data annotations [11, 12]. This review paper broadly covers 3D art such as sculptures, archaeological sites and surface art such as walls, cloth or tattoos. Additionally, it also considers 2D forms such as paintings, sketches, digital art, cartoons, logos, anime and manga. Synthetic art through the stylization of real-world data without content separation commonly suffers from the bleeding of colors from the foreground to the background and the blurring of boundaries. Despite depicting the same content, the stylistic differences between various art media highlight the importance of separating style and content in model design for different art-related tasks. To illustrate these problems, we use t-distributed Stochastic Neighbor Embedding (T-SNE), which is a commonly used nonlinear dimensionality reduction algorithm to visualize embeddings and understand similar images from the dataset according to the model, VGG-19, embedding the data. From the T-SNE visualization in Figure 1, we see that art domains of one type cluster closely, with some overlap for those sharing similar shape representations with cartoons and sketches. The large inter-class variations in the art modalities lead to their distributions looking roughly like anisotropic Gaussians, with their spread creating overlaps with images of similar subjects between the other classes. The stylization of real-world data without content separation leads to the bleeding of colors from the foreground to the background and the blurring of



Fig. 1: T-SNE visualization of the art domain as compared to the real-world images using the PACS dataset. The art modalities with paintings, cartoons and sketches showcase the clustering of art modalities close to photos, exaggerated geometries and no color or texture respectively. The dimensionality reduction uses a pre-trained VGG-19 model as a feature extractor and removes the fully connected head.

boundaries. Our paper discusses paintings, sketches, digital art, cartoons, logos, anime and manga as visual art media. It covers a variety of styles and genres such as abstract art, realism, impressionism, expressionism, surrealism, cubism, pop art. These styles range from those that resemble real world objects, leaning towards figurative art, to the other end of the spectrum using basic shapes and geometric deformations that diverge from reality.

1.1 Background

Computational art is a field that finds applications in the visual art industry for modeling collections of art [13], simulating visual art for artwork exploration [14], and replicating artworks for preservation [15]. As the field grows, it is essential to recognize that the distinction between human-created and AI-generated artwork remains subjective [1, 16, 17]. The papers discussed in the literature use datasets from the art domain for training, thus modeling the domain, even if the art dataset is a subtask or a subset of the training dataset at least. For clarity, we use the term ‘AI-generated graphics’ to refer to AI-generated works, acknowledging the ongoing debate over whether AI art qualifies as visual artwork. Additionally, when using the terms ‘art’, ‘visual art’ and ‘artwork’ we refer to the input data collection. The survey paper discusses research where AI image synthesis and generation models result in different interpretations of artworks from existing styles, diversifying visual arts datasets [18] or producing artistic counterparts to real-world images in the introduction and synthesis sections. AI tools assist in the extraction and analysis of growing collections of art to enhance interactive experiences while expanding interpretation across selected geometry priors in collections [3, 19].

It evolved from a form of static procedural art, building from preset rules with random perturbations through models such as the AARON model in the 1960s. Its later iterations were more data-centric, allowing for generation of more dynamic art using computer graphics through non-photorealistic rendering techniques that distilled artistic styles to simple parameters such as brush strokes and other learned statistics forming popular techniques such as style transfer [16]. These artworks were used to form interactive art pieces where both the artist and the audience influence the displayed design. Computational art evolved to include dynamic content generation with the evolution of image generative models, data-driven approaches learn how to synthesize AI-generated graphics through GAN-based models in the form of CAN, pix2pix, Cycle-GAN and GANVAS or learnable style transfer from DeepArt [20], digital art proliferated depictions of various media in different existing styles. Other styles involved repurposing existing computer vision models for hallucinating emergent styles from natural images such as DeepDream. With a demand for controllable AI-generated graphics with the insertion and deletion of objects of varying poses and views, the AI art community developed text-to-image models such as DALL-E and Muse. Current iterations of image generative models aim towards high-quality AI-generated graphics using diffusion models such as GLIDE or Stable Diffusion. They currently suffer from mitigating data biases from the confounding of the style and color choices of particular art movements [21] that result in hallucinating structures.

Art museums broadly use AI tools to collect user statistics regarding exhibit visitations and tours [3]. Additionally, they use computer vision in wider areas such as the Museum of Modern Art’s Thinking Machines exhibit use computational machines for artistic production[22], or the National Gallery reconstructing old master paintings using imaging techniques developed in the Art-ICT conferences. The Metropolitan Museum of Art of New York City, colloquially referred to as ‘The MET’ even sources its data, thereby helping improve the retrieval and classification performance of these tools [23]. Others use computer vision for object detection and recognition, artwork cataloging and curation, 3D tours of sculptures, and augmented reality experiences with captions describing a piece. For example, Augello et al.’s [17] cognitive architectures for artificial agents creating paintings, and other deep learning tools used in creative processes and analysis of fine art. In addition to the more passive art exhibitions, AI art has been used for enhancing the interactivity in art, such as Duan et al.’s [24] emotion analysis to create personalized art derivatives based on Van Gogh’s paintings, and Nawar’s [25] interactive art project exploring bread as a representation of one’s peculiar voice and political statement.

1.2 Geometry in the Visual Arts Industry

Identifying and analysis of artworks often use geometries ranging from global cues such as composition, perspective and proportions to identify stylistic characteristics common to an artist, to local cues such as stroke patterns, directions and shapes. Geometry is widely used in the art industry to represent perspective and lighting, as well as to reconstruct 3D shapes and locations of objects from 2D pictures. Models that incorporate proxy geometry onto artworks [20, 26] find applications in animations and VR-/AR-based museum tourism. The learned proxy geometry are geometric features or model embeddings that learn style invariances or shape and geometric data information. The 3D proxy or 3D geometric features is an intermediate representation upon which these creative applications perform operations such as relighting and novel views from their 2D projections or 2D geometric features. The use of the generation of 3D generated models extends to content recovery for art conservation projects [5] in image searching for art historians and experts. Depending on the art style the object geometry is similar or exaggerated compared to their real-world counterparts [27]. The geometry data, such as pose [28], keypoints [29] or bounding-box [30], can then be used as labels to retrieve and match images with objects that range from highly structured to highly varied geometries. They also take the form of extra input maps along with the photographs of artworks like murals or paintings on surfaces like pots, walls and robes [6] provide extra information when projected together to form 3D models. Table 1 shows examples of visual arts datasets using such geometric data for various tasks discussed over the duration of the survey paper. A 3D proxy is an intermediate representation upon which these creative applications perform operations such as relighting and novel views from their 2D projections. The use of the generation of 3D models extends to content recovery for art conservation projects [5] in image searching for art historians and experts. Depending on the art style the object geometry is similar or exaggerated compared to their real-world counterparts [27]. The geometry features, such as pose [28], keypoints [29] or bounding-box [30], can then be used as

labels to retrieve and match images with objects that range from highly structured to highly varied geometries. Extra input maps along with the photographs of artworks like murals or paintings on surfaces like pots, walls and robes [6] provide extra information when projected together to form 3D models.

Computer vision and machine learning for extraction and analysis from art collection meta data and images or 3D models provide an alternative to visual formal analysis of art collections that can be subjective, time intensive and inconsistent between experts. Geometric priors from data or those learned from models help detections in cases such as complex or cluttered scenes [11, 30, 31], exaggerated or abstract poses [32–36], and perspective distortions or low dynamic range [37, 38].

Additionally, variants of such geometric conditionals find use in generative art to create controlled and diverse outputs while preventing repetitive or biased patterns that may emerge from the underlying data or model generation process [39]. This mitigates problems such as incoherence [40], stereotypes and prejudice [41] in current AI-generated graphics that lead to texture bleeding from parts of objects to each other, deformed hands or badly constructed objects. With aspects of geometry modeled separately in the model through proxy objects, shading and illumination stages, researchers can even model impossible, inconsistent and incoherent shapes in input painting images [42] on purpose.

1.3 Paper Organization

To understand how geometry contributes to artwork tasks, we discuss the artificial intelligence techniques facilitating the use of geometry in extracting, analyzing and synthesizing artworks. We believe that discussing extraction, analysis, and synthesis is pertinent as these are the primary applications of AI tools in the art industry, utilized by experts, critics, and visitors in art collections. These processes are interconnected through the common thread of geometric considerations in various learned representations or additional constraints, which play a crucial role in enhancing the understanding and appreciation of art. We aim to emphasize the evolving nature of art and how AI tools are increasingly being integrated into creative processes and their interpretations. The focus on geometry in this discussion aims to improve the quality of the generated media by providing form guidance amidst the fluidity and variety of styles while allowing more control for artistic expression. Additionally, it addresses common failings in these AI tools [32] with regards to generated images, visual composition and geometric deformations.

We first discuss the extraction of geometric labels for humans and objects from 2D images to 3D models in **Section 2**. The object labels are divided into bounding boxes, key points and segmentation masks whereas people range from pose skeletons, landmarks and gestures. The 3D features range from explicit surfaces to implicit surfaces, and parametric models. Then, we explain the analysis of the effectiveness of the extracted geometric data on discriminative tasks in **Section 3**. The feature extraction section discusses the extraction of entities or subsets of visual art collections where geometric information is used directly as constraints or selection criteria to improve discrimination or indirectly by augmenting the collection to improve model classification. Next, we detail the synthesis and manipulation of artwork for novel

view synthesis, relighting and content restoration in **Section 4**. The section mainly discusses the use of geometry for visual art collection modification which provides new perspectives or renditions of existing artwork where any changes made are geometrically and stylistically consistent with the original artist’s vision. Additionally, it discusses content recovery where the geometric consistency is towards the original artistic medium in addition to the entity’s geometry for both global and local consistency preservation. Finally, we discuss the **future directions** for better incorporating geometry into the model architecture such that it is fully differentiable to use the full strengths of deep learning methods.

1.4 Related Surveys

Using Geometric information in artificial intelligence models facilitates the learning of representations that encode the inherent structure of visual elements. This paper covers 2D and 3D artistic visual media while discussing the extractable geometric features following it with the discriminative and generative tasks they can be incorporated into. The closest work related to ours considers geometric features at the local and image level through feature descriptors or hardware (e.g. 3D printers and scanners) for 3D models in cultural heritage [8]. Unlike our work, they focus on preservation, registration, reconstruction and enhancement, and do not consider deep learning based techniques. In this section, we first cover more recent surveys in the field before covering the works that incorporate AI and geometry in visual art.

AI-based methods have found use in creative applications such as content generation in multimedia, captioning, spoofing and AR/VR [20]. These involve models deployed for production in games such as GameGAN, storyline generation with MADE or Vid2Vid, and artwork generation from Hypercube-based NEAT that utilizes geometric regularities. A survey on computer vision in art history highlighted its applications in image search and retrieval [43] for art historians. They identify the importance of recognizing contexts such as clothing, architecture, materials, faces, patterns on objects, and artist signatures for recognizing time periods, geography and culture. Other papers focus on deep learning approaches in paintings [44] and digital art collections [45] for content recognition such as classification, retrieval and detection in images or multimodal domains. They also cover a subset of art synthesis using image generative models with losses that exaggerate styles or through latent space guidance with other models such as Contrastive Language-Image Pre-Training (CLIP). We cover the discriminative and generative tasks only if they consider geometric information in the form of annotations or pseudo geometry in the form of intermediate representations, model functions or dataset transformations to induce model invariance.

A study on mixed 2D and 3D non-photorealistic media covers the task of art conservation using multiple input spectrums [6]. They only use machine learning or statistical information for detection when dealing with the visual spectrum, with the majority of their study covering technologies for diagnosis and imaging. A similar review covers these multispectral inputs for paintings only, which extracts 3D geometric information through correspondence matching with feature descriptors such as SIFT [46]. However, we explore not only art restoration but also the repurposing

of material properties, symmetries, and corrupted regions within 3D representations. This enables us to simulate missing contents lost due to deterioration or the image sensing process.

2 Geometric Features Extraction

Object detection uses geometric features that act as descriptors learned from special classifier architectures to learn external geometry data [52] from labeled boxes or pixels that enclose the object using bounding boxes or semantic maps. These descriptors can be specialized [85] for human detection by adding more structure and robustness to affine transformations via pose skeletons at different hierarchical levels. Common geometric data capture regions of interest for both humans and objects at the semantic and instance level. The latter has not been explored in the existing works in favor of the problems of domain adaptation and limited data. Some data augmentation techniques and extra input induce style invariance and other context invariances such as time periods respectively to force models to process the input data as geometric feature embeddings that capture structure information, thereby forming geometric techniques.

Paintings are challenging to computer vision models due to their background clutter and object composition [86]. With their high diversity in poses and shapes as compared to their real-world counterparts, they can even lead to spurious detections, mistaking people for other mammals [75] or from occlusions. Furthermore, some painting datasets do not have reliable annotations with some subjects missing. To relax the problem, existing works account for the deviation of the geometry and depictions from the real world to the artistic image domain with modifications to their real-world detector training and inference pipelines.

2.1 Object-Centric Features

Object-based geometrical feature annotations involve varying amounts of information contained in excess for bounding boxes, exactly in segmentation masks and minimally in keypoints. Alternatively, data pre-processing induces the underlying class of transformations during model training or removes non-geometric information. These labels are summarized and visualized in Table 2.

2.1.1 Bounding Boxes

Bounding boxes provide the target objects' position and scale in paintings with rectangles by manual or automatic annotation using object detection techniques. Object detection is divided into single-stage or two-stage models for speed-accuracy trade-off, and scores the overlap between their proposed regions of interests and the bounding boxes [87]. In the painting domain, these models account for the gap between the artistic depiction of the object and its real-world counterpart by transforming the input data [55] or modifying its stages [11, 30, 56, 66] depending on the architecture choice. Modifications to one-stage models proceed in an end-to-end fashion to classify detected objects while multi-stage detectors optimize their constituent stages to produce better candidate regions [67].

Source	Dataset	Task Description	Dataset size	Number of classes
[47, 48]	best-artworks-of-all-time (Kaggle dataset)	Image retrieval	8.4K	50
[49]	Christian archeology (CHA)	Object detection	16K	16
[50]	StyleObject7K	Object detection	7K	10
[50]	ClipArt1K	Object detection	1K	8
[50]	Watercolor2K	Object detection	17.8K	6
[50]	Comic2K	Object detection	52.7K	6
[51]	QMUL-OpenLogo	Logo detection	27K	352
[31]	StyleCOCO	Human detection	61K	-
[52-54]	WikiArt Paintings	Object detection	81K	27
[55]	Brueghel	Object detection	1.5K	10
[56]	People-Art	Human detection	4.5K	-
[52]	IconArt-v2 (IA).	Object detection	6.5K	10
[53]	Artsy scrapped dataset	Orientation detection	2.8K	-
[57]	Pandora 18K	Style classification	18K	18
[58]	Painting-91	Style classification	4.2K	13
[59]	CMPPlaces	Scene classification	2.5M	205
[27]	DRAM	Semantic segmentation	2.5K	12
[60]	ArtSem	Semantic segmentation and Image generation	40K	5
[61]	MAFD-150	Face detection	150	29
[28]	ClassArch	Human pose estimation	1.7k	-
[62]	Artistic-Faces	Face detection	160	-
[33]	Drawings	Human pose estimation	2.5K	-
[63]	Poses of People in Art	Human pose estimation	2.4K	22
[64]	Dunhuang	Image Inpainting	5.6K	-
[65]	Chinese-Landscape-Painting-Dataset	Conditional Image generation	2.1K	-

Table 1: Artwork datasets used in the training of models discussed in the Geometric Features Extraction, Discriminative Geometric Features Analysis, and Synthesis with Geometric Features sections. The table describes the dataset names, the tasks they are utilized in along with the total size of the dataset for training, testing and validation. Finally, the number of classes in the dataset is counted, with entries left blank for single-class datasets or those using paired inputs and outputs, such as images with their ground truth semantic maps or pose skeletons.



Geometric data	Bounding box	Keypoints	2D Representation Object segmentation	Landmarks	Pose skeleton
Existing work	[11, 12, 30, 51, 55, 56, 66–70] [49, 61, 71–77]	[63, 73, 76, 78]	[11, 27, 72, 79]	[62, 80, 80–82]	[32–36, 74] [81, 83, 84]
Techniques used	Single stage detectors Multi stage detectors Transfer learning Multimodal learning	Correspondence matching Point-source detection	Segmentation models Conditional maps Saliency maps	Off-the-shelf face detection Landmark transformation Geometric style transfer Region networks	Transfer learning on pose estimators Multimodal pre-training Pose matching and clustering
Visualization					

Table 2: Different types of geometric labels extracted for objects and humans ranging from bounding boxes indicating scale and position, points of interest indicating correspondences, maps and masks for important regions, and landmarks and pose skeletons indicating structure. The table lists the techniques used to extract these geometric labels in the Extraction section for images of objects and humans. The bounding boxes better localize people in backgrounds with higher contrast for the impressionist painting of two children sitting with a piano. The keypoint detection here detects pose as the points of interest in a portrait image of a lady. The segmentation task provides shape information through the mask along with a bounding box to enclose the extent of the boundary for the people on the pier and the boats. In landmark detection, the structure of the face of a Hindu God is annotated with points and the contour connecting them.

Single-stage detectors such as You Only Look Once (YOLO) directly classify objects and perform regression using Convolutional Neural Networks (CNN) to get their location and size from the predicted box coordinates and object aspect ratio. They utilize spatial consistency in datasets to implicitly embed the geometric and content information together and do not use intermediate features to represent object regions. In fine-art paintings, associations between objects provide a non-destructive means for identifying visual connections for investigations into its history or authenticity [67]. However, datasets with nonnatural collections lack strong spatial correspondences among highly diverse objects, limiting the effectiveness of pre-trained detectors. Furthermore, the shape abstractions in paintings and sketches vary drastically from photographs in the resultant high-level feature space [68]. To address these challenges, detectors often employ data augmentation techniques like style transfer [51, 55] to align real-world images with artistic image styles and structures, thereby increasing the artistic images dataset size. However, these augmentation techniques require further processing to maintain semantic consistency [18] since failed samples actively detriment model performance. These augmentations may compromise semantic consistency [55] unless complemented by techniques such as mask content mixing, which combines valid regions of objects from multiple images without overlap or obfuscation [51]. Nevertheless, style transfer is useful for texture-biased models like CNN-based object detectors [55], unlike shape-biased models [24]. For example, style transfer is unsuitable for certain artistic image styles such as stick figures and sketches [56] where careful content mixing that preserves semantic correspondences is preferred instead [88].

Multi-stage detectors add an additional feature extraction stage before classification and regression which customizes the learned object representation, size and shape. It is commonly based on Region-Based Convolutional Neural Networks (R-CNNs) that have a feature extractor to extract and warp the regions of interest into a uniform aspect ratio with a CNN for model prediction [87]. In paintings, the learned deep features help detect near duplicate objects that only differ in style [67] that pay homage to popular artists and schools by copying or modifying their composition. This ability enables tasks, like grouping styles or achieving texture, invariance by replacing [30] or modifying architecture stages [11]. Earlier methods extract static features using style templates [30, 69] to capture the appearance variation of an object so that the feature extractor works with a robust representation that accounts for a small variation of style, color, scale and orientation. The template-based detection finds use in objects with standard configurations such as paintings authentication which requires artist signature attribution from small motifs in architecture and heterogeneous paintings of still life, portraits and landscapes. Some later work improves parts of pre-trained R-CNNs with multi-head attention modules [11] for improved model performance to extract and exaggerate the most class discriminative subregions to improve detection for cases in the Art-DL dataset’s rarer classes Dominic and Paul or those that co-occur such as Mary and Jesus. The selection stage in detectors helps separate the foreground from the background in crowded scenes composed of objects of varying sizes, unlike single-stage detectors. Furthermore, unlike single-stage detectors, the modifications to the feature extractor allow generalization to object geometry in terms of scale and

proportions, but they limit the adaptability to novel tasks based on the choice of the feature extraction strategy. Depending on the method of feature representation learning, detectors can enhance their performance for classes with small sample sizes, which are often considered anomalous in accordance with their training dataset. This process involves separating the inputs’ appearances from their other attributes [89, 90]. In cases where the classes are not imbalanced, the representations can be robust enough to handle label shifts or produce consistent activations with specific groups [91, 92].

Transfer learning involves training the detector’s pre-trained layers on a different data domain, using additional components and data augmentation to adapt the model for a new task. The performance of a pre-trained model relies on how the second stage of multistage detectors is trained. These training choices may involve training parts of the models in stages to bias the detector toward the desired task. Alternatively, bootstrapping the model can provide weak supervision by using a small subset of clean, labeled images to propagate to the remaining data [70]. For Japanese Ukiyo-e paintings, although pre-trained models accurately detect faces, the model’s classification of the cropped face bounding boxes is poor without fully tuning the model weights [93, 94]. Fully fine-tuning classifiers can significantly improve model performance, especially when coupled with data-efficient learning methods like contrastive learning [94]. However, this approach depends on the sampling strategy for the training data, and the compatibility of data augmentation with the pre-training dataset and task at hand.

Instead of retraining or modifying a model for a specific task, we can leverage information from multiple domains and utilize relationships in the latent space between multiple models for domain adaptation with multi-stage detectors. For instance, multitopic language modeling [66] can provide context for image captioning by using scene relationship embeddings from the painting description dataset SemArt alongside pre-trained Faster R-CNN and Residual Network (ResNet) models. This approach requires fewer annotations for generating image descriptions, but the domain gap among the models results in significant variation in image description evaluation metrics. In some image captioning tasks [66, 71], the object detector acts as either a feature extractor or a caption generator. Features from different modalities’ outputs are fused using the attention mechanism of transformers. The detectors play dual roles as image captioners and intermediate feature extractors, facilitating semantic alignment at either the image-text [66] or image-image [71] level, thereby bridging different domains. In the former case, the detector can also categorize detected objects into a hierarchy of textual categories. Moreover, semantic metadata can enhance object detection performance [72] by filtering out objects incompatible with the periods depicted in scenes, eliminating the need for auxiliary models.

2.1.2 Patch-Based Region Selection

In the absence of labelled bounding boxes, patch-based selection methods approximate object locations with feature engineering and simple classifiers. They enforce faithful geometry through spatial correspondences [67] embedded in model features with a further selection stage that accounts for a particular class of transformations. These matches at the feature level narrow down the area of interest from the image to a patch level to return the target object. Object retrieval for paintings typically uses a

multistage model for feature engineering and better detection and localization using part-based models. These multistage detectors are useful for noisy datasets where the target classes are absent in the object. The first stage clusters related image-level candidates through text mining and choosing patch-level candidates through segmentation and MLDPs through a Histogram of Gradients (HoG) while the second stage trains class-wise DPMs as a detector.

During the feature engineering stage, the fusion of different data modalities provides the model with a shared representative space within the same domain [73] or between different domains [66] and robustness to noise from different sources. Additionally, their feature embeddings often require heavy pre-processing with dimensionality reduction into embeddings such as fisher vector [76], histograms and templates to encode image content. Other common embeddings involve model-specific image embeddings such as HoG as a learned feature encoding method [74] for a middle-level discriminative patch as a robust template matching alternative after hyperparameter tuning.

For detection, part-based models work with a smaller number of proposals and utilize the patch statistics alongside coarse geometric information from the preprocessed features. Deformable Part Model (DPM) [74] is a popular choice as a class-specific sliding window detector [12, 75] since it is robust to object detection under a lot of arrangements of its sub-parts. Detector robustness is crucial for domain adaptation in problems such as detecting gods or animals in vases [12] where there is a lot of subject ambiguity. Siamese networks as the detector architecture provide another solution by contextualizing the data as a strategy to capture the scene with co-occurring object embeddings [49].

2.1.3 Geometric Transformations for Content Selection

Simple geometric transformations like affine transformations and cropping [95] also help models train on small datasets while learning invariance properties from the data.

Data augmentation techniques provide the additional benefit by mitigating the need for extensive data labeling by generating additional training examples that capture variations in object appearance and their contexts. There are techniques for generating new data from the image level to preserve the training distribution. LogoMix [51] synthesizes samples from overlapping different logo patches, effectively preventing the model from overfitting to synthetic data and ensuring it learns robust features from real-world data.

Style transfer creates a hybrid image that preserves the structure of the content image of the real-world while its style takes after the painting images. Other data augmentation techniques often drastically transform the training distribution to account for the gap in diversity between domains [96]. They can have fidelity towards different aspects of the input pair depending on if they were trained on multiple styles, a single style, or through iterative optimization. The synthetic data aggressively transforms the training dataset, provided the input data is not independent of texture such as edge maps while needing only a limited set of style images.

Style transfer also suffers in diversifying images with structures [31] like stick figures, but it can still provide robustness to textural distortions to adapt object detection in the painting domain. Their classifier uses ResNet-152 and Faster-RCNN that are fine-tuned on stylized COCO dataset and showcases the improvement in detection on larger training datasets, even if they are simply stylized versions. While style transfer provides a large shift in the color and texture information, without trading off the content loss fidelity towards the style image, it does not correspondingly warp shapes. In object detection, style transfer helps reduce the cross depiction problem to that of color and structure [77] attributes in the image while ignoring geometric information that is context dependent like gesture, shape and pose. Artistic representations can be more shape-biased [31] for stronger geometric features and better model performance. The Computer Vision Group (CVG) system [77] performs image retrieval using contours from images for stroke information and negative example training after expert annotation with five bounding boxes to determine spatial extent and geometric relation.

2.2 Human-Centric Features

Human detection on the other hand involves highly regular structures at the face, body and hand level captured by bounding boxes, pose skeletons and landmarks. Pose and body shape information provides information for perception and identities [35] through the proportion of their parts. Human poses can vary in depiction across time periods and represented with different topology or motifs [63] where the referential deviations can represent artistic signatures or movements. They can also increase the difficulty in detecting poses by blurring contours, distorting proportions or occluding joints through apparel or other objects or lighting, sometimes even changing the cardinality of the parts to indicate mythological creatures. Face detection in uncontrolled settings like modern artistic styles [61] is very challenging due to occlusion as well as variations in shape, color, texture and face size. In end-to-end training, the ratio of the preservation of the geometric features or the pose regression loss to the style transfer loss is vital, unlike training them as separate parts for achieving high pose accuracy [34]. There is still a cross-domain generalization problem since the joint positioning is more accurate with real-world images compared to artistic images, but more data from stylization gives better performance after a cutoff. Paintings retrieval can benefit from pose annotations at different levels of abstractions [83] to compensate for mislabeling. Inverted label propagation produces these levels of poses through producing annotations induced from the source to the target image provided that the dataset size is sufficiently large.

2.2.1 Hand Gestures

The detection of hand poses involves the position and orientation of the hands and their fingers with respect to the body in the form of templates or pose skeletons. In portraits and paintings, they commonly form hand signs and iconographic meaning with irregular finger positions or unnatural gestures with hand actions [97]. They indicated a group's, family or religious memberships and ranks, personality traits

and an artist’s signature style. Learnable hand templates [69] strongly separated the hand from the background, which convolve with Laplacian of Gaussian filters across the image and give strong responses on contour alignment. While their collections capture the primary variation of appearances across scale and rotations with data augmentations and principle component projections, they do not encode relations between the hand and body. This results in detecting false positives for speaking gestures and other semantically ambiguous actions.

Follow-up works found that the use of deep pose estimators alone results in poorer accuracies in western fine-art paintings [37, 38] due to distortions of perspective and low contrast of the body against the background. They used the OpenPose model [98], a multistage convolutional model to detect and match part-wise confidence and affinity maps, to detect the skeletal pose of the hand with 21 keypoints and if it is left or right-handed. The model learns body part relations, their locations and orientations with their corresponding confidences with learned heatmaps and vector maps. The detection improves with better representations learned by pose descriptors. For example, the ResNet-50 pre-finetuned on a large sign language dataset can effectively recognizes gestures involving both hands when compared with their simple angle pose key point descriptor [37], but fails on the less represented classes with hand-object interactions. Such misclassifications also result from low interclass variations in gestures in which similar poses belong to different classes.

2.2.2 Facial Landmarks

Facial landmarks in paintings have more variations compared to their real-world counterparts leading to its architectures disentangling the style and pose into separate multistage models or a need for data augmentation.

Multistage models such as the Cascaded Pose Transform network (CPTNetV2) can simulate head and face pose animation [80] to model pose displacements while inpainting the facial features separately to disentangle the problem into the two separate poses transformations. These models need a refinement stage to add details while maintaining consistency. Then, a fusion generator utilizes both the pose information that was disentangled by imposing masks to guide their individual generations. Other two-stage models can detect modalities like bounding boxes and keypoints for the human figure [81] from the photograph domain to that of paintings using a semi-supervised learning method with transformers through a teacher-student model distillation. It predicts a fixed set of proposals for each image, removing the need to account for overlapping boxes and imbalance between the foreground and background. Distilling geometric information for domain adaptation provides better results than fine-tuning or style transfer with additional label conditioning.

Artistic augmentation[62] for landmark detection requires image transformation through style transfer techniques followed by a part-based feature correction step for landmark warping to account for structural shifts and decorrelating parts. Techniques like part-based correction and tuning and Geometric style transfer account for extreme styles and higher variation in landmark points. The stylized portraits’ landmarks are warped to the mean facial shape vector of the target style to capture a signature

structure using Thin Plate Spline (TPS) interpolation, but the method cannot handle fundamental shape variations from natural faces in anime or manga portrayals. A ResNet encoder-decoder and region networks can account for global and local landmark arrangements [82] and result in accurate prediction of inner facial features in high-resolution images. Style transfer and geometric augmentations, to randomly shift or resize facial landmarks along with their movements on a TPS displacement field, account for their changed arrangements in artistic faces. Despite suffering from jawline landmark localization due to ambiguous labeling, the method works well for salient portrait features for the eyes, nose and mouth.

2.2.3 Body Skeleton

Poses between people in paintings and photographs [32] can be effectively aligned by pose detections and matching them through geometric transformations. The latter validates and measures similarity under different scales and positioning while making the detection robust to noise and missing parts. The method is not robust to unknown poses, occlusions, ambiguous poses, or any spurious connections that arise from these challenges. Style transfer can bridge the gap between photographs and paintings for both person and pose detection in curvilinear surfaces like vases and create a dataset to fine-tune the HRNet [99] model for the tasks. With a perceptual loss on both tasks, the model can adapt the annotations to the pose and detection losses with the stylized data.

To compensate for limited painting data, the pose estimators can also be pre-trained from 3D renderings [33] of artistic media like anime or manga which provides joint positions from the underlying rigging. These models can simply be fine-tuned on a smaller dataset of drawings to effectively ignore the problem of domain gap from models pre-trained on photographs. The pose information can be utilized in other tasks such as image retrieval. Pose similarity followed by clustering helps retrieve similar paintings [36] through methods like K-medians with metrics which are invariant to scaling, rotations and translations after detecting them through pre-trained models like OpenPose.

2.3 Segmentation Masks

Image segmentation partitions the image into pixels that group into multiple classes which can be further grouped into individual objects that belong to the same class in the case of instance segmentation. This requires fine localization of objects in the scene regardless of scale, occlusions from clutter or other objects, or appearance changes from lighting or environmental conditions. Earlier works use deformable models, which provide an object shape template representing a distribution of warped objects, and graph cut to partition an image into regions while providing boundary separation [100]. By rephrasing the image segmentation as a deformable model optimization problem, they represent the Chinese paintings by their unique color choice in neighborhoods and the direction of texture. The deformation model splits the image into connected regions, while the texture directions represent flexible sparse foreground/background features like edge convolution filters to detect orientations.

Deep learning based segmentation networks like DeepLab v3, a fully convolutional model segmenting objects at various scales with spatial pyramid modules and cascading dilated and upsampling convolutions, transfer well to the artistic domain with transfer learning and style transfer [79]. When trained only on natural images, the baseline model on modern human portraits outputs faulty segmentations due to weak lighting cues, different color and texture choices compared to photorealistic images, and similar contours in the object and its background (e.g. striped sleeves and sofa). The model improves when fine-tuning once on style-transferred images, before training it on the real portrait images. Despite the unnatural color tone transfer in some regions, the model only fails with extreme shadows, reduced style, flat regions and unnatural skin tones.

Similarly, style transfer augmentation [27, 101] improves the result of segmentation models by increasing their shape-bias while simultaneously providing robustness against various image corruptions (e.g. noise, blur, adverse weather conditions such as fog, motion blur from fast-moving objects). By separating the task into coarse binary mask proposals and fine mask refinements [88], the segmentation model becomes robust to domain shifts. The detector, without any fine-tuning, is robust enough to find objects of various styles from watercolor, clip-arts to comics. The binary mask extracts multiple objects that belong to the foreground or the background using cosine similarity measures on features from models like self DIsillation with NO labels (DINO), a self-supervised transformer that learns object semantics from global representations forming from local image patches. The detector with their novel loss refines these masks and adds undetected regions from the mask proposal step.

3D Representation	Explicit	Implicit	Parametric
Existing Works	[102–105]	[106–111]	[112–115]
Sub Types	Point clouds Voxel grids Mesh structures	Neural Radiance Field Gaussian Splatting Signed Distance Field	Skinned Multi-Person Linear Model
Visualization			

Table 3: The extraction of 3D geometric data in the form of parametric representations, explicit and implicit surfaces. The table mentions the papers under these categories along with the representation types used in the subsections under 3D Features. The explicit model is a triangular mesh of sculpture and its wireframe representation. The implicit model is a Gaussian splatting rendering from a single view using the DreamGaussian architecture. Finally, the SMPL visualization is a rendering from a Contrastive Language-Image Pre-Training (CLIP) based text prompt of an Asian mermaid.

Going forward, more advanced input and output features could enhance segmentation performance. Multimodal image input features [73] could be extended beyond artistic object segmentation, as they derive from various data sources and feature extractions, and compensate for missing information across modalities. Additionally, they address issues like low-contrast cracks in photographs, absent cracks in IRP, and noise in X-ray images. Incorporating pseudo segmentation maps as an output feature has proven effective, leveraging discriminative features via Class Activation Maps. This approach not only improves image localization and detection simultaneously in paintings [11] but also holds potential for adaptation in human segmentation.

2.4 3D Features

The feature representation of 3D models can take generic forms such as implicit neural structures and explicit geometric structures or specific versions from parametric constrained deformable models. Implicit neural networks represent 3D shapes and surfaces in a pointwise manner with their learned function, while Explicit models represent the objects as a 3D point collection and differ in usage by the efficiency and ease of use of their data structure for different tasks. Parametric models, on the other hand, use a small number of shape and pose parameters to efficiently represent objects in a class like Skinned Multi-Person Linear (SMPL) for human bodies with strong constraints to prevent large structural deviations. These representations are summarized with visualizations in Table 3.

2.4.1 Implicit Models

Implicit geometry structures can separate the natural scene geometry from artistic stylization by utilizing a two-stage model with a Neural Radiance Field (NeRF) and a 2D stylization decoder which gets the projected view to style [106]. The desired style can be customized by conditioning the latent code that is the input to the decoder which also serves to deblur the rendered scene from the NeRF depth output. The model requires multiple stages [107] for view projections and style transfer to mix their outputs to get a stylized scene. When transferring style onto a mesh, unseen style inputs result in blurry reconstructions with naïve methods transferring only the overall color tone of the style image. By rendering the geometry and stylistic aspects separately, the ARF paper [108] showcases the transfer of the subtle textural details of the watercolor feather image onto their Family statue scene example.

To optimize the transfer process [108] for better representations in view extraction, the image level style transfer uses a deferred back-propagation at a patch level to accumulate over all the patches at the neural field. To reduce computational complexity, the final model can only consider the components in the visible field of view [109] but still requires all direction illumination and material information. Alternatively, better representations can be learned using an image generative model like a Style-GAN [110] and fed to the NeRF module, facilitating detailed feature extraction while conditioning the generator on geometric priors such as pose. For specific types of extractions such as human sculptures, implicit models like Pixel-Aligned Implicit

Function (PIFu) that are trained on human data produce better results while accounting for domain shifts by adapting the intermediate features with a Maximum Mean Discrepancy (MMD) loss that aligns their moments and handles topology shifts [111].

2.4.2 Mesh Structures

Direct 3D representative models such as voxel grids [102] ignore object artifacts found in the neural field representations in two-stage scene stylization. These are limited by the resolution of the extracted 3D model and the computational size of the intermediate features. Perspective changes warp the projected image from moving the lines of convergence constraints [103], thereby changing the vanishing point. They achieve this view emphasis by warping the quad mesh and the corresponding homographic matrices while constraining the projection geometry.

It is possible to extract more faithful 3D models using a dictionary of surface gradients and exploiting the symmetry of such mesh structures in other views. The utilization of such self-similarity with inpainting in point clouds [104] finds applications in reconstructing damaged and structurally deformed architecture and sculptures.

Stylistic renderings of multiple views for stereo paintings can use 3D paint strokes on top of partial grid mesh structures as a two-stage model [105].

2.4.3 Parametric Models

Artistic 3D models such as sculptures closely resemble the human figure, but are limited in dataset size or consist of larger variations in pose or structure to provide emphasis to foreground or background characters, sometimes exaggerating the shape from certain viewpoints. The SMPL model provides a minimal, resilient pose and shape representation through low-dimensional vectors. The customizability of intrinsic parameters accounts for the anatomical differences in artistic statues, making it handy for reconstructing statues like the Wounded Amazon with a different sized arm or for those with missing limbs such as the Esquiline Venus [116]. By adding a Signed Distance Field (SDF) to check the occupancy of particle effect selected by the artist, the body can be textured and locally manipulated by keeping track of the normals and tangents on the SMPL’s deformable mesh [114]. Instead of manual deformation and texturing, the SMPL model can be simply extended with a CLIP loss to make the mesh representation similar to that of text control with a differentiable renderer [115]. Integrating these joint interactions and their confidence coefficients with transformers increases the accuracy of human reconstruction and the speed of the mesh extraction [112]. The model works well when the character images are clear, but failing in abstract paintings in Picasso’s works where parts of the head are missing or incomplete. When the body parts are occluded, it leverages joint relations to recover the skeleton topology and pose but encounters significant errors, particularly with parts like feet that have fewer adjacent joints. They use a High-Resolution Net (HRNet) to extract these human features. Representative keypoints from all of its output keypoints (e.g. nose representing the face) are then selected and the model fuses the joint and mesh information using a graph transformer model. In 3D scene extraction, template skeletons can be conditioned with bas-relief geometry, contours and silhouette information, for particular styles of sculpting to estimate 3D skeletal poses from 2D poses [113]. The

Method	Dataset	Task	Metric: Value	Source
MMD and PiFU	ScanTheWorld scrapped meshes	3d model extraction	Chamfer Distance: 0.047	[117]
Styled Deeplabv3	Neural Style Transfer on the Baidu People Segmentation dataset	Segmentation	IoU: 74.9%	[79]
Multi-style feature fusion	LTLL	Object detection	Accuracy: 90.9%	[30]
Styled HRNet	ClassArch	Human pose estimation	mAP: 49.4%	[28]

Table 4: This table represents methods in the surveyed geometric feature and data extraction papers, showcasing their datasets, tasks, highest performance measure score, and sources.

choice of the 3D mesh, such as B-mesh that provides good deformation for animation and edge flow, is separated from the rest of the scene while keeping distance information to jointly model trees, animals, and environment cues like drapery that are commonly found in these bas-relief sculptures.

2.5 Effectiveness of Geometry-based Methods in Extraction

Object-centric tasks benefit from input region proposal selection strategies and additional geometric labels to add context to the input or act as pseudo-ground labels. Region selection by voting helps find and localize small motifs, achieving a maximum retrieval performance of 91.3 mAP for the LTLL dataset where other models suffer from selecting regions with low correspondences [30]. Mixing of regions selected without overlap and missing content helps in data augmentation for small datasets for logo detection for an improvement of 7.05% mAP [51]. Encoding the context instead of missing regions from cropping helps improve fine-tuning object detectors by 3.5% mAP for unseen categories with an additional 2.5% for seen categories [49]. Leveraging this prior knowledge of pre-trained models and semantically aligning them with a newer domain helps improve performance even with difficult subdomains such as abstract paintings [71]. Geometrically enhanced annotations also enhance the quality of the training dataset with maps enhancing salient regions [73] while eliminating irrelevant areas in crack detection or time-specific label predictions for object detection in paintings [72].

Human-centric labels such as facial landmarks and hand or body pose differ in depictions in the paintings domain, resulting in poor results with simply fine-tuning the model on the task. Without accounting for style, body pose estimators reach less than 60% AP [75] while face detection algorithms reach less than 35% F1 score on modern face datasets [61]. When stylizing images without modifying the corresponding pose, models can get an improvement of 6% on the mAP even without labeled data [28]. Style-tuned models that pre-train on stylized content and poses gain an improvement of 36.7 mAP for the specific task of pose estimation and 34.5 mAP for the more generic person detection task. In image retrieval, performing geometric verification after fast annotation matching retrieves a longer sequence of visually similar links as compared to other models that simply match feature embeddings.

Parametric model-based extraction methods utilize prior knowledge to account for lower-quality data, modeling complex environments with missing information, and

leveraging geometric information in the modeled latent space. They provide additional benefits in terms of a reduced computation time due to the strong prior with FAKIR [116] extracting each iteration of a modeled statue in 9s. Additionally, it provides precise joint positions and bone radii with better shoulder location estimates compared to its counterparts, thereby producing geometrically consistent artistic 3D models. JointMETRO [112] also achieves painted sculpture reconstruction despite occluded human poses for incomplete models by utilizing this prior knowledge of human body joints. While the parametric model provides an alternative to explicit ground labels, other techniques such as domain adaptation can build upon them to reconstruct higher-quality sculptures with a Chamber Distance of 0.04 [117].

3 Discriminative Geometric Features Analysis

Various painting analysis tasks utilize geometric features from low level (*i.e.*, *local feature descriptors such as brushstrokes or optical flow maps encoding direction*) and intermediate level (*i.e.*, *cross spatial correspondences between objects to identify keypoints and landmarks*). These analysis methods even utilize outputs from feature extraction methods such as a list of bounding boxes. These geometric features and data are utilized in the tasks that are sub-categorized into scene classification, retrieval and style classification.

Scene classification relies on similarity measures to determine object arrangement and assign scene labels. It involves three core stages: feature extraction (e.g ResNet without fully connected layers), spatial correspondence encoding (e.g Attention, K-means clustering), and output mapping (e.g. Aggregation of bounding boxes, multiclass classifiers with Softmax activations). In contrast, style classification focuses on identifying unique styles in artistic images, which can vary in visual cues rather than content. Notably, stylistic manipulations in images can impact feature extraction and geometry due to object detectors, posing challenges on both scene classification and retrieval. Scene retrieval aims to find images resembling a reference by mapping output features to identify closely matching candidate images of a specific scene class.

3.1 Object Detection

Objects in paintings have large shape exaggerations in modalities such as cartoons, vary drastically in their composition with cluttered scenes consisting of objects of different scales and spatial arrangements. The spatial layouts are crucial in scene understanding [11, 49], with objects representing visual motifs for artists or indicating time periods and culture by their co-occurrence with other objects. These artistic datasets are small in size [50], with some providing only image-level annotations [118, 119] or missing object-level labels [11, 49].

Traditionally, object detectors in landscape scenes focused on analyzing and understanding low-level features, such as brushstrokes, for capturing scene dynamics using optical flow although the result can be suboptimal when noisy object regions are not effectively detected as principal components. By integrating these low-level details with region-based segmentation algorithms like Comaniciu’s mean shift clustering, which groups input scenes by color, the system provided a deeper understanding of

the composition at the object-level. The segmentation information encoded by the clustering method contains the region information distribution according to the object thereby capturing the variations in appearances and their relationship with each other through. This data-driven approach allowed for more refined scene interpretation when fed into a threshold-based classifier, facilitating a clearer distinction between objects and background elements [119]. More recent computer vision techniques used models pre-trained on a larger domain for the same classes in a target domain by fusing the style from the artistic target image modality and the content of the source modality. To create the synthetic pair, methods like arbitrary style transfer using Ada-IN, which learns a style transformation network to translate images from one domain to another [118]. Such a method provided easier access from faster training and no fine-tuning to multiple sub-tasks (when considering multiple modalities in domain adaptation) is needed, unlike learning generative models such as GANs. The work shares the pre-trained backbone and fully connected classifiers with multinomial logistic losses from a domain confusion loss to predict the domain of the image and an object classification loss. Multiple modalities force the network to learn a general representation, enforcing style invariance with the choice of style transfer affecting the retained structure and details of the synthetic image. Realistic paintings get limited improvement while modalities that emphasize shapes with their contours like cartoons and sketches gain substantial performance gains.

When the feature alignment process between domains uses a generative model such as Cycle-GAN [120] instead of style transfer, it loosens the requirement of pairwise source and target domains for image translation. It also provides a fully differentiable domain adaption method where the multiscale detector, Retina-Net, acts as another discriminator for multiple adversarial losses, one for domain confusion and the other for object prediction to understand the variations in artistic content across domains while constraining the learned transformation to produce an object of the target style. In galleries, these models are utilized to study how variations in lighting conditions, viewpoints, and mixed environments affect the model’s ability to correctly interpret and adapt to painting or sculpture regions. The interaction between artistic images or 3D models and real-world surroundings creates complex and varied input data, and analyzing these variations helps reveal strengths and weaknesses in the model’s alignment process. While the model excels at translating artistic styles and maintaining target domain features, it struggles with scenes that involve multiple artistic images ranging from paintings, clip-art and comics, where the co-occurrence of certain object classes can mislead the interpretation of the discriminative regions. Effective analysis of detected objects depends on the quality of feature alignment from the source and target domains, particularly when the dataset comes from varied artistic image modalities [50]. By modeling parallel object proposal networks, the classifier can better handle variations in data from fusing regions and adjusting its parameters based on the distribution through XGBoost. The boosting algorithm helps emphasize more difficult or rare cases, which is crucial for understanding less frequent objects or features within the dataset. However, despite the model’s strong performance in multi-scale

analysis and fast inference provided by YoloV5, it often struggles with datasets heavily focused on people, where searching for other, (e.g. less common) objects becomes more challenging without further modifications to the pipeline.

3.2 Style Classification

Style classification involves artist identification and the common visual elements, techniques and forms used in their works. The artists attribute the forms to lower-level textures, such as their choice of color palette, brushstroke, or materials, up to the higher-level choice of fine-art painting compositions. Style classifiers benefit from feature fusion techniques that merge geometric image representations with deep learning features as input to a multiclass classifier. These geometric representations are handcrafted for the problem to account for the large inter-class variation in styles and class imbalances stemming from artist-based classification, which result from variations in the artists' prolificity. In older works, CNNs, were used purely as object feature extractors which is less effective in capturing image representations compared to learned ensembles of handcrafted features like Classemes or PiCoDes [54]. However, their performance improved significantly when the object region was first extracted and used as input to the model. This approach highlights the importance of isolating relevant regions for analysis, enabling CNNs to better understand and represent the essential characteristics of the image, thus offering a more accurate interpretation. When a DPM provides class-specific regions to a multiscale CNN to provide a holistic encoding and learn a distribution of local encodings through a GMM, their joint embedding after aggregation through techniques like Fisher vector gives better performance [58]. More recent work has enhanced CNNs' ability to analyze and understand image representations by incorporating discriminative signals from an SVM-based classifier. This approach refines the clustering criteria, allowing the model to generate centroids that more closely align with the original target label distribution [57]. This combination of deep learning and SVM-based analysis facilitates a deeper understanding of the underlying data structure, ensuring that the representations captured by CNNs better reflect the true characteristics of the target labels.

3.3 Scene Classification

Scene classification involves the general subject matter or the semantics in the painting and considers categories like outdoor and indoor-based scenes, landscapes and portraits, seascapes and landscapes, still life or other labels describing the scene type. Simple methods like segmentation (eg. Normal Cuts) can extract visual descriptors like HOG or GIST from regions at the image level [121] for representing the structure and texture within each segmented region. These object descriptors are used as the input to a Bayes classifier to form the RoI pooling operation in a multistage object detection model. Due to the simplicity, the classifier focuses on colors and results in mistaking images from nudes and portraits since they both contain skin colors or they can not distinguish between cityscapes and landscapes due to the latter being the superset. Later models use CNNs to represent generalizable feature representations

Method	Dataset	Task	Metric: Value	Source
Faster R-CNN with CAM and ResNet-50 backbone	ArtDL 2.0	Object Detection	mAP: 41.5%	[11]
MLCNN	Artsy, WikiArt paintings	Orientation classification	Accuracy: 92.42%	[53]
DPM detector	painting-91	Style classification	Accuracy: 74.8%	[58]
GMM	CMPlaces	Scene retrieval	mAP:14.2%	[59]

Table 5: This table represents methods in the surveyed discriminative analysis papers, showcasing their datasets, tasks, highest performance measure score, and sources.

for multiple modalities alongside constraints like MMD that force a shared representation among different CNN heads [59]. With modality-specific fine-tuning, the target dataset can have a smaller number of samples while the distribution constraint enables an emergent alignment of objects shared across multiple representations.

3.4 Human Perception Analysis

Objects in paintings can appear distorted despite being portrayed with the correct geometry based on the viewer’s vantage point from large visual angles that tilt and straighten, reduced saliency of peripheral objects, to depth-wise elongations [122]. Experiments involving participants to move towards the painting until they saw the object of interest take the desired shape or subtend an angle showcased similar results to that of projective geometric analysis, but to varying extents. When measuring their response, the farther vantage points had varying perceptions of distance in peripheral objects for large paintings or those approximating 3D scenes. A case study of Piranesi’s painting composition hints at possible approaches to balance the trade-off between accurate scene geometry against perceptive distortion [123] with the pieces utilizing projections from multiple viewpoints along the central vantage line. While this reconstruction shows inconsistency using geometric restitution as a tool for analysis, it does provide a view to different proportions, sizes and relative distances of non-distorted objects from different viewing angles.

3.5 Effectiveness of Geometry-based Methods in Analysis

In the task of cross-domain object detection, additional representations such as CAM [11] or context encoding [49] can compensate for a lack of ground truth (GT) label while bringing in contextual information as learned by pre-trained models on a larger, well-annotated dataset. CAM acts as a pseudo-GT beating the SOTA on the weakly supervised object detection task in the IconArt dataset by 14%, while with the context encoding the finetuned model beat the SOTA in the mean average precision (mAP) by about 3.5% at 0.25 intersection over union (IoU) for UnSeen categories.

For the broader task of scene and style classification, the content of the painting has the greatest effect on the classification accuracy favoring methods that process the local and global regions separately before integrating the results. For example, the number of classes from both painting categories and the number of directions can be reduced to an umbrella, holistic classes to increase the average accuracy to 90% for orientation classification [53]. In the latter task, part-based models or models that

extract text and style separately improve model accuracies by 6.4% and 13% respectively, compared to the non-contextual cases. This approach allows a more detailed understanding of the distinct elements within the data, enabling the model to better capture and interpret the contextual relationships between features, and thus leading to more accurate predictions. In other cases, the feature representation can be conditioned to include context such as meta-data [124] for an improvement of 26% as compared to a 6% increase from building upon context-aware solutions.

4 Synthesis with Geometric Features

The synthesis section covers the generation and manipulation of artistic images or 3D models ranging from paintings, cartoons, sketches to sculptures. The generative models include components to separate style from structure, including geometric deformation modules or networks along with a separate module to blend the separated components together. Part of the model store or use the geometric feature as input to a different part of the generative pipeline. Geometric data in the form of masks and semantic maps, 3D representations such as point clouds or polygon meshes, are used as additional input to these models. In image manipulation, some regions of the image are changed by adding and removing objects or to match the deformations as in a reference image. The task of novel view synthesis covers unseen artistic 2D/3D data sampling using image generation models, relighting to provide a view of the image with different lighting, time lapses or seasons, and rendering to change the local geometric details. These main techniques are summarized with visualizations in Table 6.

The synthesis section covers the generation and manipulation of images or 3D models of artistic data ranging from paintings, cartoons, sketches to sculptures.

4.1 Image Manipulation

Image manipulation involves deforming the contents of the image to that of a reference image or editing the objects into or out of the source image. The objects immersed in the scene can belong to the same or different artistic modality while the region deleting the object uses the neighbouring area to fill in candidate outputs. The main challenges during the manipulation of artistic images are artifacts around the boundary or bad correspondence matches indicating a semantic gap.

4.1.1 Style Transfer

Style transfer offers the benefit of considering the holistic geometry as compared to traditional image processing techniques at the cost of smooth and artifact-free stylizations. Artistic 2D and 3D data can benefit from this separation since it often geometrically deforms images as a stylistic choice. It can be formulated as optimization or transformation of the stylization and geometric modules at the pre-processing, model-level and post-processing stage.

During style transfer pre-processing, we learn the distribution of shapes [125] while promoting textural invariance using augmentation to improve object detector performance by tackling implicit model bias. However, there is overlap through style



Synthesis Task	Style Transfer	Inpainting	Relighting	Conditional Image Generation
Existing Papers	[34, 125–128, 128–138]	[64, 104, 139–141]	[14, 142, 142–150]	[60, 126, 151, 151–156]
Types of techniques	Geometric priors Model induction via perceptual losses Post-processing and style refinement	Mask-based editing Structure completion	Color transfer Portrait lighting	Generative Adversarial Network Variational Autoencoder Diffusion model
Inputs	 source  reference	 source  source	 source  pose	
Outputs				

Table 6: The table lists image synthesis methods for image stylization, manipulation, relighting and generation with their inputs and outputs. It also lists the techniques used for these tasks with their papers in the Synthesis section. Style transfer involves the deformation of the flower vase still life images in reference to that of the source. Inpainting on the other hand fills in the missing region indicated by the mask in the impressionist-style source image with that of the reference still life image. It results in the source image of a lady near a dining table replacing the mask with a similar flower vase from the reference image, but one matching the blue color tones of the source. Relighting changes the foreground and background lighting of a source image from a cavalry riding through a meadow in spring to winter. The Image generation task is visualized here as conditionally regenerating a sketch in the pose of the reference image. It changes the man preaching with a book to one with spread arms.







Image Domain	Content Image	Style Image	Neural Style Transfer	Stylized Image Adaptive Instance Normalization	Deformable Style Transfer
Same Domain					
Cross-Domain					

Table 7: Visualization of stylized outputs from different style transfer techniques. The style and content images are either from the same 2D/3D artistic data domain or the content is from real-world images to represent the case of using style transfer as purely a data augmentation technique. These style transfer techniques are covered in the Synthesis section and show the difference in outputs from the simple iterative style transfer method to one that uses the style transformation network and finally, one that is geometrically aware. The style of the abstract depiction of the lady’s face is muted for the first, while the correspondence between the facial features is preserved in the 2nd style transfer output. When the geometric deformations of the style are also taken into consideration by matching correspondences, the source images of a painting or photo of a person match the abstractness of the style.

per-training to projection or shapes and geometric augmentations from deformations or distortions of size and orientation. We can pre-process style transfer in the extracted mask for the structural warp [126], but we need to refine the mask boundaries for better texture transfers. We could also embed class-specific warp feature fields through style content image pairs and their corresponding vector field [127]. These warp fields account for intra-domain and cross-domain generalization. Alternatively, we can trade off the mesh style to structure integrity by weighting the topological optimization against the style loss [128]. These geometric deviations include unconnected objects or unfaithfulness to the overall main design, requiring post-processing filters for imposing global geometric features.

The methods that use a specialized model for handling geometry in tandem with style transfer require preconditioning or pretraining and cannot handle multiple representations. Stroke-based rendering can model strokes and the image content separately with a neural painter as an image generative model and style transfer’s content loss respectively [129]. These strokes can be approximated as a joint fusion problem in the Fourier domain [130] by shifting the brush stroke from the source image patch to that of the mean patch. Vector fields along with noise can also model brush stroke stylizations of the input image [131] while preserving object-level details with their contours untouched. In facial transfer, geometric flow models [132, 133] with facial landmarks blend and warp details together with an adversarial or deformation loss model the addition and subtraction of the attribute while preserving the person’s identity.

By using style transfer as a post-processing task, it can be treated as a separate optimization task on the output of an existing pipeline [34, 134]. This extends its application to input and output images in the 2D and 3D domains.

To preserve geometric features in 2D images, we need three images to provide the style and geometry information of the stylization, as well as the target geometry or matching algorithms to provide the correspondences at the image or model-level. Geometric Style Transfer [135] uses a geometric warping network with a specialized style image alongside loss pairs for structure and style. When transferring style to objects detected in bounding boxes, geometry-aware style transfer can fail if there are no good semantic matches between the style and content images [136]. Neural style transfer can synthesize designs with a mesh reference and a novel topology optimization inspired compliance loss [128] to separate the geometric information from texture. Bounding boxes facilitate the warping of specific matches of regions of the content and style images [136] before applying style transfer for the desired texture while retaining multiple object structures. Alternatively, the selected regions can constrain the style loss, providing spatial and aesthetic effects, error correction and controlled transference of styles according to the masked regions [157]. Thus, the structural loss can be modified with warping, reconstruction or geometric consistency terms to achieve geometric deformations and projections that match the style images.

When translating an image to a 3D model, stylized outputs could preserve their pose after style transfer [34] using a novel style reconstruction loss on the 2D stylized image which is lifted to the 3D depth image as a pose regression task. Arbitrary style transfer allows the stylistic transfer of the pose only, where the style loss adds a new HSV colorspace loss for color insensitive style transfer along with a slightly

modified style and content loss to follow the general direction of the style and content features than a reconstruction loss. The added self-supervision loss corrects the lifted stylized pose in the 3D space when reconstructing the stylized image, thereby using the geometric nature of the bone map with style transfer.

3D stylistic modeling uses a separate stylistic module or loss to transfer the style while preserving the geometric features in an implicit or explicit representation. By sketching the contours at an angle of a sparse point cloud [137], the artist can interactively reimagine the generated 3D model from the initial models of pre-existing categories, with contours retrieved from viewpoint matrices. With the help of the pre-existing dataset, the user input contours can be matched and retrieved for that part of the viewed model. Similarly, we can represent the structure of 3D point cloud vases as features extracted [138] to transfer shapes with style transfer losses and a laplacian loss to preserve local details such as edges, contours and patterns better.

4.1.2 Inpainting

Inpainting in artistic 2D/3D data structures is used in applications for image [139] and region level[104] completion, editing[141] and restoration[64]. They do not preserve part of the content in the reference image but transform the style of the target object to that of the source. They perform image completion in selected patches or masks.

Structure completion in Inpainting incorporates the target boundary-based patch selection as a search and voting optimization task [158]. This patch matching incorporates geometric transformations such as reflections, non-uniform scaling and other perspectives through affine transform approximation. Alternatively, region filling is possible from data fusion from overlapping local patches with gradient-based self-similarity [104]. Using cropped patches, we can achieve image completion, which produces ringing artifacts [139] with vanilla image generators like Generative Adversarial Networks (GANs) that follow a loss of high-frequency information. These methods help remove boundary artifacts within the patches but not on the patch edges themselves.

Mask-based inpainting commonly incorporates explicit geometric features such as object masks from scene segmentations [79]. Their flaws are mainly from the properties not included in the mask types themselves such as scene segmentations with depth discontinuities and changing boundaries from ghost pixels [100] that are mistakenly attributed to different objects. Inpainting specific aspects such as high-frequency details of a painting is possible through edge maps[64] in a different domain or colorspace. These masked images allow editing of subjects in the erased area with generative models like latent diffusion models and Large Language Model based prompt guidance [140, 141].

4.1.3 Conditional Image Generation

Generative Adversarial Networks help in the quick sampling of outputs with multiple controllable attributes embedded as a conditional vector, an additional input that commonly represents a style or shape vector. Its attributes are formulated as disentangled representations to form independent control factors, with multiple modalities (such as

other input domains for artistic modalities or geometric labels) encoded with Cycle-GAN variants [60, 151] to facilitate unpaired domain-wise translations and learning their correspondences. Variants of Generative Adversarial Networks can be used for synthesizing warped stylization images, for example, using facial landmarks or edges from line-art [159] as the conditioning vector in conditional Cycle-GAN. Conditioning inputs such as segmentation maps provide size and location cues, however, the model produces worse results if the semantic classes and their object appearances diverge significantly from the pretraining data distribution. On the other hand, keypoints indicate local relationships and correspondences, which help generate results with significant shape exaggeration. However, the mismatches or deformations can generate implausible results that do not match the input distribution. The choice of conditional for GAN style guidance influences the geometric embedding network’s ability to capture finer, localized details like strokes through directional fields, thereby shaping the diversity of synthesized styles. [152]. By borrowing the stylistic deformations from the source image and the reference’s colors and content, the model generates samples that belong to a new synthesized AI-graphics movement with its real-world natural input distribution. In place of external conditional inputs, the correspondences between parts of the paintings can be approximated by cycle consistency losses while additional losses such as brush stroke and ink wash losses help the model simulate geometric textures such as brush strokes or the washed-out effects of paints such as Chinese inks [160].

Diffusion models [153] extend the level of control to different shapes and styles while also learning attributes that adapt the conditional’s data distribution. These models learn an iterative mapping from a simple distribution, such as the Standard Normal, to a complex distribution. In the case of latent diffusion models, they learn the latent space’s distribution to better represent and mix the input distributions. The encoder of the geometric labels [126] such as sketches and segmentation maps are frozen and the outputs are fused with the input image embeddings. For the former, the temporal order of sketches can be encoded into another latent space with part-wise Autoencoders or Variational Autoencoders to reduce computation [154, 155]. Compared to GAN-based methods, they can generate better results with people with accessories and facial features, as well as supporting different levels of part abstractions. Conditioning on segmentation maps [156] provides coarse geometry that can be enhanced with other labels such as text that make them adaptable instance labelled maps. While these pseudo-labels can be weighted to control their influence in the result, the properties of the geometric label are not shared with the text descriptions without clearly defined instance boundaries. Text labels serve to offer global and semantic context clues within spatial layouts, yet they do not directly enable manipulation of the layout itself [161]. This helps to clarify and rectify ambiguous and erroneous learned correlations that may arise from coarser maps. Since the encoders do not share information, the embeddings cannot learn the shape separations in the mask encoder which extends to other embeddings of geometric cues.

4.2 Novel View Synthesis

Novel view synthesis in artistic image and 3D model datasets refers to techniques that change the perspective of the scene and its objects by relighting and recolorization.

This changes the focus of the subject in the scene and aligns the illumination process to be closer to real-world settings to better convey the artist’s intentions. In these artistic images, the illumination sources can drastically vary from realistic sources in terms of color, direction, and intensity which encourages modulating the lighting on the extracted geometric shape of the painting.

4.2.1 Relighting

The manipulation of directional lighting in paintings [14] allows art historians to gain new interpretations of the artistic images by determining the nature and effectiveness of optical instruments in the past. Adjusting the lighting to account for cast shadows, specular reflections, and self-shadows, while incorporating point source illumination information, can reveal previously unnoticed elements of the composition. By examining the direction of illumination, new insights can emerge, such as identifying geometric inconsistencies or uncovering occluding contours, which may hint at image tampering or offer alternative realist painting compositions.

Simple techniques like illumination template matching [142] help in lighting transfer from the source to the target image by warping the matched face descriptor if the light source is simplistic and there is only one subject. In interior lighting, the light sources from the style image can be transferred to the content by extracting the perspective information with key-point detection to warp the surface map according to the style elements [143] before restoring the perspective. Similarly, face illumination descriptors on an active shape model along with deformation transformations can perform facial illumination transfer [142]. Utilizing 3D models as shading proxies allows for lighting transfer from the user-provided light source direction [144] or segmented reference object [145] to the target shading proxy. This proxy contains artistic style, brush strokes and color information to learn an implicit normal and depth map, which overcomes the limitations of previous methods on handling highly stylistic scenes. Alternatively, the use of CLIP [146] for obtaining physical lighting properties and local geometry information can transfer lighting using explicit normals and materials. Similarly, a 3D mesh structure can be stylized according to the text prompt with CLIP guidance [147] while keeping the differentiable rendering of the correct viewpoint and lighting of the final 3D mesh.

4.2.2 Recolorization for Artistic Time Lapse

Artistic Time Lapse decomposes the painting into its constituent objects with the help of frames of the video creating the artistic animation to relight them according to the environment’s illumination source for a day or across seasons. The methods detect keyframes with color shifts to get the art’s decomposition into layers indicating depth from the artist’s perspective and predict the next frame by learning the sequencing using a Conditional Variational Autoencoder [148] that conditions on the previous frame. To generate time lapses of the painting itself instead of the painting process, the albedo map is estimated for each frame, clustered with their linear layering to pick the colors not affected by lighting and hue shifted with the artist’s choice of colors to produce the time-lapse effect [149]. The decomposition process can be differentiable to extend the process to other materials and allow the artist to control the levels of

lighting across the layers [150]. These methods do not consider a data-driven approach due to the lack of availability of a digital art dataset with layer information which can encode the relationships between layers for more complex lighting.

4.3 Content Recovery

Damaged artistic mediums such as paintings or sculptures can be recovered by remodeling them as 3D structures and in-painting the missing regions from noise and occlusions during the imaging process or from material wear and tear. The deformations in the imaged data can occur at the surface or subsurface level in the case of paintings. Depending on the nature of the artistic image or 3D model, the data acquisition method varies from Photogrammetry for paintings and generating structure from motion with LiDAR sensors for sculptures.

4.3.1 Remodeling

Art conservation helps in the objective diagnostic and documentation using photogrammetric remodeling tools [162] to annotate parts while allowing comparison study with its representation and the acquired data. Furthermore, it helps in the objective diagnostic and documentation using photogrammetric remodeling tools [162] to annotate parts while allowing comparison study with its representation and the acquired data. When modeling rock paintings while considering ease of extraction, details preservation and accuracy of the reconstruction for the non-domain experts [163], the following characteristics of 3D point clouds were observed. They generated highly detailed models at short distances to the camera with a large number of redundant photos, regardless of the camera quality.

Many 3D reconstruction techniques require the 3D modellers and art critics, historians or curators to work in tandem with the extracted model [164, 165] with regular geometry to add in missing context or clean the rendered structure. The quality of the reconstruction depends on the retained details while trading off the computational expense of the generated representation. The former can utilize inpainting using self-similarity in point clouds [104] to help in reconstructing damaged and structurally deformed architecture and sculptures. It allows the creation of more faithful 3D models using a dictionary of surface gradients and exploiting the symmetry created from artistic intent in other views of such mesh structures. If the paintings have subjects with geometric features close to real-world people, the images can be used directly with off-the-shelf models such as 3DME [166] for 3D face reconstruction from 2d egocentric portraits. The latter uses modifications on the 3D reconstructed point cloud to reduce their structure computation, with more points needed for ornate structures as compared to flat surfaces. The point cloud is compressed using a multi-resolution Octree (3D representations using tree data structures) and converted to a polygon mesh with a photographic texture overlayed upon it, with the overlaying requiring extensive pre-processing for cleaning the rendered model. The 3D reconstruction gives problems in perspective inconsistencies and changes in the pose without explicit geometric guidance [164], some of which are unnatural due to artistic liberty.

4.3.2 Painting Medium Surfaces

Western paintings on surfaces with various shapes were previously studied in detail [5] at different scales and curvatures, forming murals, frescoes and pottery decorations depicting different numbers of subjects in each piece. Macro-level objects are highly variable in condition, nature and global shape, and thus typically contain a single subject. Data collection is highly dependent on the acquisition strategy, retrieval and characterization algorithms for collections such as statues and pottery. Curvilinear paintings require their canvas to be collinear with the surface curvature to prevent distortion and stretching [4]. The folds and warps along with the tensile strength of the canvas are necessary to model the deformation instead.

4.3.3 Subsurfaces

Painting restoration is interested in the surface topology of paintings and it requires information such as material, subsurface, layers or deformations. It involves data acquisition from non-intrusive scanning and stitching the overlapping spectral scans [167] before detecting cracks and restoring the artistic image [168]. The crack patterns can be simulated for authenticating paintings [169] for the behavior of the outer film, the interaction between their layers and shrinkage from drying using a physics-based system. Any inconsistencies between the paint layers stress amongst each other and elasticity from an estimated time period could imply painting fabrications. These misalignments can be highlighted as a learning or visualization tool [164] by superimposing the elements in the reconstruction with that of the paintings. During 3D printing of oil or other substrate paintings [170], the reproduced image is susceptible to staircasing effects due to the thickness of the layering of colors that differ in ink density and viscosity. The paintings must maintain accurate geometries and color information which is obtained through a Point Cloud after matching and decoding images taken from different orientations. Using their novel color layering order, they account for the curvatures of the pigments and boundaries of their overlaps to reduce the staircase effect while simultaneously maintaining the quality of the color reproduction. The identification and restoration of damaged paintings utilizes nondestructive testing methods like infrared thermograms that measure heat emissions, over time from absorption and emissions, and over the pixels in the image. These use Multistage models that focus separately on the temporal aspect using an MLP which returns patches for spatial processing using a U-Net [171] segmentation model to generate the reconstructed image.

Multispectral imaging opens new interpretative possibilities in painting restoration and 3D surveying by exploring material features while capturing spectral data that are often missed with traditional imaging techniques such as X-ray radiography and infrared reflectography. Furthermore, they reveal hidden layers, underdrawings, and pigment compositions that are invisible to the naked eye [172]. Geometric deformation analysis using close-range photogrammetric techniques can evaluate deformations such as craquelure patterns, color raisings, detachments or engravings in the range of $\pm 0.1\text{mm}$ with more advanced equipment putting the range in $50\mu\text{m}$ or sensor triangulation [173].

Novel Evaluation Metric	Paper	Description	Usage	Task
Faithfulness score	[132]	Distance of two cropped facial landmarked regions of a source and target attribute on a normalized feature space	The lower the better for a more faithful transfer	Style Transfer
Semantically Corresponding PSNR (SC-PSNR)	[174]	PSNR on the MSE of fixed patches surrounding the corresponding key-points of two images to provide a patch-level ground truth measure.	Higher the better for similar matches.	Novel View Synthesis

Table 8: The novel evaluation metrics as used for the task of AI-graphics synthesis (Section 4.4). It lists the metrics introduced in the cited papers before their definition and usage description. Finally, it mentions the task for which each metric is used.

4.4 Evaluation for Synthesis Methods

The main evaluation methods for AI-graphics synthesis tasks can be divided into 3 types:

- User studies to quantify human perception
- Quality measures through deep learning models or at the pixel level
- Performance measures through loss terms and statistical measurements

Table 8 indicates the novel evaluation measures introduced in these papers to evaluate regions of AI-generated graphics focused around geometric labels. Finally, the section includes a discussion on the effectiveness of incorporating geometry into various synthesis tasks.

4.4.1 User studies

To study the quality of the generated results, existing works conduct user studies of different group sizes, with larger studies conducted with the Amazon Mechanical Turk (AMT) [148, 175, 176]. The smaller studies range from visual assessments, comparison studies and the likeliness to the human creation processes. To evaluate the image quality and the aesthetics of the AI-generated graphics, they employ Likert scales [107, 108, 177, 178] or mean opinion score tests [179] of different ranges with 3 [180] as the lowest and 10 as the highest [136]. However, they could be used as a comparison study to compare the researcher’s models with baselines [132, 181, 182] under multiple criteria, such as realism and consistency in style or the coherence of attributes like text. Visual Turing tests [65] form a branch of these comparative tests, where the users have to judge whether the artistic output data is human-made or machine-made along with their level of confidence in their answer. The AMT tests find use in reaching a broader audience [176] and designing studies to help sequential stroke-based models identify salient regions in generating AI-graphics and mimicking the human creation process [148, 175]. Finally, user studies can evaluate the ease of use and the versatility

of the proposed tool, such as in tasks like relighting objects and people in different environments[145].

4.4.2 Quality Measures

Quality metrics are those that measure salience, image quality in pixel or feature space and semantic relevance by alignment of different modal features. We consider the salience measures that specifically incorporate geometric information to evaluate the importance of regions or pixels. The metrics that involve explicit geometric labels are faithfulness score [132] and SC-PSNR [174]. The former uses facial landmarks and the latter employs keypoints, as detailed in Table 8. Others evaluate the geometric labels estimations like SOA score [178] that uses an object detector. It follows with other metrics like IoU for segmentation maps [183], foreground L2 distance with a pre-trained Deeplab-v2 model [184] and shape quality through MSE on estimated depth maps and poses [148, 185]. However, Semantic relevance measures change their evaluation process by using instance segmentation to mask out the image background and crop to the object [186]. This allows for using image quality measures like PSNR, SSIM and LPIPS without background interference [64, 107, 185]. These geometric evaluation metrics have distinct advantages over image quality measures. For example, metrics like SSIM are less sensitive to color changes and are susceptible to blurring and low contrast, similar to PSNR [187].

4.4.3 Performance Measures

Performance measures are used to evaluate other properties of the model beyond data quality, such as diversity, identity and task accuracy. To model diversity, some works use average gradients [185] to consider both clarity and texture variation in the image while influencing the content detail. Some papers use accuracy for preserving the identity [110] and semantic relevance [108, 178] of the detected regions, which mitigates inconsistent results from the identity shift with better results from applying the metric to both the geometric label prediction and the AI-generated graphics. Moreover, models can generate finer geometric labels with better shape quality through metrics such as IoU [183, 184] for similarity matches and fine-tuning coarse mask predictions. Additionally, loss metrics such as L2 with Nearest Best Buddies (NBB) correspondences help closely reconstruct keypoints [136]. Other error measures include MAE, RMSE and Classification error [80] to add constraints for class and pose consistency.

4.4.4 Effectiveness of Geometry-based Methods in Synthesis

Geometry can be utilized in many ways for the task of AI-generation of graphics, ranging from geometric-style embeddings, separate geometry and style latent spaces, geometric annotations to condition the generator, and refining geometric annotations to form fine geometry conditionals. The first case shows improvements of 10% compared to baselines without geometric reference [128] in the case of topology transfer, but has little constraint over local to global geometry consistency. The second case shows varying results in diversity of geometry and visual quality as in the case of MW-GAN [106] with improvements from the worst-case geometry generations but a better

average visual score. The third case produces outputs similar to Chinese landscape paintings in the case of SAPGAN with an agreement of 55% compared to baseline GANs with 11%. Finally, fine-tuning the geometry of the conditional provides generators with better data to start training with an average improvement of 3.7 IoU across its datasets for unsupervised segmentation generation.

Qualitative experiments on these methods indicate aspects such as color and detail of the generated results that these geometry-aided models excel in. In image restoration, geometry guidance with features such as edges improves color coherence and sharpness in local regions. Other models overlap regions undefined by the stroke or pose models leading to ghosting effects [179] or blurry regions [185]. The generated regions sometimes fail to preserve the appearance, detail and color information [80, 146] in the case of extreme poses for the task of anime head animation. Conditional maps also force the model to reduce uncolored regions or in the opposite case, leaving out topological features such as legs in the case of sculptures [111]. They are particularly useful in 3D AI graphics generation, where detailed meshes help with ambiguous geometry, view consistency or artifacts [187]. More free-flowing AI-generated graphics models, on the other hand, could use elastic regularization or coarser geometry to help with under-constrained problem domains. Poor correspondences from mismatching geometric annotations result in unwanted behavior in the output like spatial discontinuities, overlapping objects [136] and their misalignment [188]. The deformation and blending operations of objects in the final rendering partly attribute to this malformation.

5 Limitations

The papers discussed in the extraction, analysis, and synthesis sections frequently introduce new datasets as part of the technique’s novelty, which complicates benchmarking against existing state-of-the-art (SOTA) techniques and diminishes the reliability of evaluating their effectiveness. Many of the cited papers in the synthesis section lack datasets with ground truth for image or 3D model manipulation tasks, instead employing techniques such as neural style transfer to bridge the domain gap. Consequently, the geometry of synthetic datasets becomes more aligned with real-world data, but they lack the stylistic composition choices inherent in artistic image and 3D model collections. Some other papers train and test on web-scraped collections that lack curation.

One significant limitation of the synthesis section is the lack of comprehensive experiments evaluating various models on different artistic modality datasets. While we explored multiple synthesis techniques, our evaluation was restricted to a limited set of models that were mostly ablation studies comparing model components and capacities. This constraint hampers our ability to fully understand the comparative strengths and weaknesses of each technique across different scenarios. Additionally, the use of a narrow range of datasets limits the generalizability of our findings. To provide a more robust and conclusive analysis, future work should include extensive experiments with a broader selection of models and diverse datasets of artistic images, including those from various geographic regions and demographics, with a particular focus on sculptures and 3D models. Additionally, implementing a standardized performance

metric that evaluates the quality of AI-generated graphics quality beyond qualitative experiments, is crucial, since measures like PSNR, SSIM, FID miss image semantics.

6 Future Directions

The papers in the literature point towards promising future directions by exploiting better similarities in data and data or model-level transformations. We cover the following four main directions:

- **Improving Data Quality:** To utilize human-in-the-loop annotations for better data tailored to a problem by learning the refinement of model predictions using experts.
- **Addressing Domain Gaps:** To learn correspondences between two domains to account for domain gaps while preserving semantic regions as determined by the geometric conditional.
- **Fine-tuning Geometric Controls:** To control the behavior of the conditioning input on the output to provide soft, learnable levels of constraints on the trade-off between style and content in the final stylized output.
- **Conditional Geometric Labels:** To use conditional geometric labels to constrain the latent space to allow model simplification by abstractions across multiple levels.

6.1 Automatic Data Annotations

The encoding of interactive annotations allows the model to learn and generalize artists annotating data. This provides an advantage over differentiable augmentation techniques by providing targeted annotation as compared to changes in the augmentation to maximize the model performance not completing the geometric label. Off-the-shelf models provide an initial annotation of artistic 2D and 3D data that are missing geometric label information. These can include the interactive annotation UI with which the corrected annotation or its editing click controls [189, 190] is encoded as the conditional with the input. A refinement stage or module corrects the annotation to go from coarse to fine labelling, filling in missing or undetected structures. Alternatively, we obtain accurate object annotations by refining the model prediction by alternating between the model learning stage and the human annotator correction stage[191]. Thus, learnable labelling provides an added benefit of resolving overlapping predictions due to multiple input sources with the expert curating the predicted labels. Finally, we can build upon these concepts to form a fully automated annotation model from SAM [192], an object segmentation model that produces high-quality masks for both real-world and artistic images. By using model-predicted annotations with their user-corrected versions, it forms semi-automated annotations which the model can finetune upon to generate reliable, fully automated masks.

6.2 Attention-Based Cross-correspondences

The domain gap between real and non-photorealistic images depends on the learned correspondences between modalities, whether they are from different artistic mediums or between text descriptions and images. The attention mechanism mixes up features between the multiple input domains, resulting in a shared representation bridging the

inputs’ domain gap [193]. The mismatches between the domains result in extractions of poor geometric labels such as segmentation maps [194] that fail in fine-grained alignment of classes. Although the learned alignment produces noisy results that make the training process unstable, the refinement of the outputs mitigates the issue by correcting these false positives. Using the cross-attention mechanism, the structure of the resultant object depends on the interaction between the conditional and image embedding [195]. Injecting cross-attention maps across the model layers can bias the model towards spatial layout and geometric relationships, such as spatial co-occurrences derived from learned feature correspondences.

6.3 Controlled Guidance

Existing multi-attribute guided AI-graphics generation pre-define attributes to incorporate in the resultant images, and keep the controls independent. The choice of the guidance mechanism in diffusion models allows control in the influence of multiple attributes in the geometric labels or text prompts in the resultant image. Controlling the influence of geometric labels on the output lets the model bias outcomes toward input domains or indirectly from the disentangled content and style representations. Fine-tuning the model with learnable adaptors that transform geometric annotations to embeddings provides structure information to the resultant output [196] while preserving the generator’s learned latent distribution representing the input. These adaptors represent individual attributes with their weights needing manual adjustments to combine into the desired output. They also lack consistency between the local changes to the global view. With energy-based generation models, the conditional information can influence the intermediate steps of the generation while keeping the artistic images dataset distribution intact [197]. The geometry labels can affect the trajectory of the latent space interpolation and get removed at the output sampling step by treating it as the stochastic term that vanishes at the end of the Stochastic Gradient Langevin Dynamics model. This allows the model to select the contribution of the conditioning labels and their constituent parts towards the output.

6.4 Geometric-Aware Models with Object Embeddings

Flow-based and multistage hybrid models incorporate geometric information to represent the input globally, which does not allow for granular control of geometric details. These use geometry information to redistribute salient regions, but do not implicitly allow the articulated parts as multiple attributes to control generations. The correspondences of the embeddings of the disentangled geometric labels and learned structure information in the generator models result in complex shapes [198]. These require heavily annotated labels to condition the inputs, but allow the models to learn the relations between the attributes while allowing the label to define the level of granularity in a controllable generation. The latent geometry provides additional benefits such as transferability to non-photorealistic images preserving the overall shape and pose of semantically similar subjects, but produces less details and robust generations if the labels lack descriptiveness [199]. Additionally, any learned appearance variations are constrained to the part regions, thereby providing local control. This learned joint

correspondence commonly fails with highly stylistic inputs with inconsistent views and drastic poses [200] which are common in more abstract painting styles.

7 Conclusion

This review delves into the effects of leveraging geometric data within deep learning architectures for artistic tasks for extraction, analysis and synthesis. During extraction, we examine the choice of models and their transferability to the artistic images, 3D models and animation domains. This evaluation handles rough versus fine data annotations, ranging from region-based detectors of multiple stages to fine-tuning with data augmentation. During analysis, we consider their performance in discriminative tasks according to the granularity of the available geometric labels. Models with coarser geometric labels typically require task-specific visual features that encode spatial relationships, while those with finer annotations generally need fewer additional modules to achieve higher accuracy. In finer geometry, Region-based object detectors outperform parts-based models in large-scale datasets but have poorer outcomes in abstract paintings or occlusion-heavy datasets. Moreover, the review highlights the effectiveness of deep learning models like DeepPose, OpenPose, and Mask-RCNN in detecting fine geometric labels compared to traditional computer vision methods, despite limitations in detecting atypical poses or occlusions in special cases. The significance of 3D geometric data surfaces is demonstrated by their advantages in accurately capturing object shapes and providing depth information, employing volumetric meshes, implicit functions, and parametric models. Furthermore, it outlines their utilization in comparing poses between statues and paintings, exploiting pose tracking, and leveraging artistic medium attributes. The extracted geometric data has a wide variety of applications in digital artwork analysis, ranging from image retrieval and scene classification to style identification and semantic relationship identification, are thoroughly examined. This highlights the effectiveness of deep learning techniques and geometric priors in scene classification and object identification. Lastly, the paper underscores the pivotal role of geometry in synthesis and manipulation tasks within computer vision, showcasing its contribution to maintaining object geometry, stylizing novel views, and enhancing image details without color bleeding or loss of local information. Additionally, it sheds light on how shape constraints and conditioning priors facilitate image refinement, super-resolution, and digital art conservation, harnessing geometric data to guide brush strokes, enhance image details, and simulate 3D models in conservation efforts.

Acknowledgments

This project is supported in part by the EPSRC NorthFutures project (Ref: EP/X031012/1).

Declarations

The authors declare that they have no conflict of interest.

References

- [1] Bellaiche, L., Shahi, R., Turpin, M.H., Ragnhildstveit, A., Sprockett, S., Barr, N., Christensen, A., Seli, P.: Humans versus ai: whether and why we prefer human-created compared to ai-created artwork. *Cognitive Research: Principles and Implications* **8**(1), 42 (2023)
- [2] Hirsch, A.J., Stocker, G., Jandl, M.: *The practice of art and ai*. Hatje Cantz Verlag (2021)
- [3] Rani, S., Jining, D., Shah, D., Xaba, S., Singh, P.R.: Exploring the potential of artificial intelligence and computing technologies in art museums. In: *ITM Web of Conferences*, vol. 53 (2023). EDP Sciences
- [4] Sklodowski, M., Pawlowski, P., Górecka, K.: Geometrical models of old curvilinear paintings. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)* **8671**, 578–585 (2014) https://doi.org/10.1007/978-3-319-11331-9_69
- [5] Pintus, R., Pal, K., Yang, Y., Weyrich, T., Gobbetti, E., Rushmeier, H.: A survey of geometric analysis in cultural heritage. *Computer Graphics Forum* **35**, 4–31 (2016) <https://doi.org/10.1111/cgf.12668>
- [6] Borg, B., Dunn, M., Ang, A., Villis, C.: The application of state-of-the-art technologies to support artwork conservation: Literature review. *Journal of Cultural Heritage* **44**, 239–259 (2020)
- [7] Remondino, F., Rizzi, A., Barazzetti, L., Scaioni, M., Fassi, F., Brumana, R., Pelagotti, A.: Review of Geometric and Radiometric Analyses of Paintings. <https://doi.org/10.1111/j.1477-9730.2011.00664.x>
- [8] Pintus, R., Pal, K., Yang, Y., Weyrich, T., Gobbetti, E., Rushmeier, H.: A survey of geometric analysis in cultural heritage. In: *Computer Graphics Forum*, vol. 35, pp. 4–31 (2016). Wiley Online Library
- [9] Redmon, J., Divvala, S., Girshick, R., Farhadi, A.: You only look once: Unified, real-time object detection. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 779–788 (2016)
- [10] Mathieu, A., Inria, T.M.U., Russell, B.C., Aubry, M., Sivic, J.: Painting-to-3d model alignment via discriminative visual elements. *ACM Trans. Graph* **33** (2014) <https://doi.org/10.1145/2591009>
- [11] Milani, F., Vago, N.O.P., Fraternali, P.: Proposals generation for weakly supervised object detection in artwork images. *Journal of Imaging* **8** (2022) <https://doi.org/10.3390/jimaging8080215>

- [12] Crowley, E.J., Zisserman, A.: Of gods and goats: Weakly supervised learning of figurative art. *learning* **8**, 14 (2013)
- [13] Özgün, F.N.K., Alaçam, S.: A computational approach for analysis of art compositions. *Gestão & Tecnologia de Projetos* **18**(2), 109–121 (2023)
- [14] Stork, D.G.: Mathematical foundations for quantifying shape, shading, and cast shadows in realist master drawings and paintings, vol. 6315, p. 63150 (2006). <https://doi.org/10.1117/12.681141>
- [15] Li, Q., Zou, Q., Ma, D., Wang, Q., Wang, S.: Dating ancient paintings of mogao grottoes using deeply learnt visual codes. *Science China Information Sciences* **61**, 1–14 (2018)
- [16] Hertzmann, A.: Can computers create art? In: *Arts*, vol. 7, p. 18 (2018). MDPI
- [17] Augello, A., Infantino, I., Manfré, A., Pilato, G., Vella, F.: Analyzing and discussing primary creative traits of a robotic artist. *Biologically Inspired Cognitive Architectures* **17**, 22–31 (2016) <https://doi.org/10.1016/j.bica.2016.07.006>
- [18] Ernst, H.: *Artificial: A study on the use of artificial intelligence in art* (2023)
- [19] Fan, X., Liang, Y.: The research on the characteristics of ai application in art field and its value. In: *4th International Conference on Language, Art and Cultural Exchange (ICLACE 2023)*, pp. 146–160 (2023). Atlantis Press
- [20] Anantrasirichai, N., Bull, D.: Artificial intelligence in the creative industries: a review. *Artificial intelligence review*, 1–68 (2022)
- [21] Srinivasan, R., Uchino, K.: Biases in generative art: A causal look from the lens of art history. In: *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency*, pp. 41–51 (2021)
- [22] James, B.: *Thinking machines: Art and design in the computer age, 1959–1989, the museum of modern art, new york, usa, november 13, 2017–april 8, 2018*. *Design and Culture* **10**(2), 219–223 (2018)
- [23] Ypsilantis, N.-A., Garcia, N., Han, G., Ibrahimi, S., Van Noord, N., Tolias, G.: The met dataset: Instance-level recognition for artworks. In: *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track (Round 2)* (2021)
- [24] Duan, Y., Zhang, J., Gu, X.: A novel paradigm to design personalized derived images of art paintings using an intelligent emotional analysis model. *Frontiers in Psychology* **12** (2021) <https://doi.org/10.3389/fpsyg.2021.713545>
- [25] Nawar, H.: Collective bread diaries: cultural identities in an artificial intelligence framework. *AI & SOCIETY* **35**, 409–416 (2020)

- [26] Cox, B.D., Berns, R.S.: Imaging artwork in a studio environment for computer graphics rendering. In: *Measuring, Modeling, and Reproducing Material Appearance 2015*, vol. 9398, p. 939803 (2015). <https://doi.org/10.1117/12.2083388>
- [27] Cohen, N., Newman, Y., Shamir, A.: Semantic segmentation in art paintings. In: *Computer Graphics Forum*, vol. 41, pp. 261–275 (2022). Wiley Online Library
- [28] Madhu, P., Villar-Corrales, A., Kosti, R., Bendschus, T., Reinhardt, C., Bell, P., Maier, A., Christlein, V.: Enhancing human pose estimation in ancient vase paintings via perceptually-grounded style transfer learning. *ACM Journal on Computing and Cultural Heritage* **16**, 1–17 (2022)
- [29] Lorente, O., Riera, I., Chaudhuri, S., Catalan, O., Casales, V.: Museum painting retrieval. arXiv preprint arXiv:2105.04891 (2021)
- [30] Ufer, N., Lang, S., Ommer, B.: Object retrieval and localization in large art collections using deep multi-style feature fusion and iterative voting. In: *Computer Vision–ECCV 2020 Workshops: Glasgow, UK, August 23–28, 2020, Proceedings, Part II 16*, pp. 159–176 (2020). Springer
- [31] Kadish, D., Risi, S., Lovlie, A.S.: Improving object detection in art images using only style transfer. In: *2021 International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8 (2021). IEEE
- [32] Jenicek, T., Chum, O.: Linking art through human poses, pp. 1338–1345 (2019). <https://doi.org/10.1109/ICDAR.2019.00216>
- [33] Khungurn, P., Chou, D.: Pose estimation of anime/manga characters: A case for synthetic data (2016) <https://doi.org/10.1145/3011549.3011552>
- [34] Wan, Q., Lu, O.: Napa: Neural art human pose amplifier. arXiv preprint arXiv:2012.08501 (2020)
- [35] Islam, M.T., Nahiduzzaman, K.M., Why, Y.P., Ashraf, G.: Informed character pose and proportion design. *The Visual Computer* **27**, 251–261 (2011)
- [36] Marsocci, V., Lastilla, L., Pozo, S.D., Kainz, W.: Geo-information pose-id-on-a novel framework for artwork pose clustering (2021) <https://doi.org/10.3390/ijgi10040257>
- [37] Bernasconi, V., Cetinić, E., Impett, L.: A computational approach to hand pose recognition in early modern paintings. *Journal of Imaging* **9**(6), 120 (2023)
- [38] Bernasconi, V.: Gab - gestures for artworks browsing. In: *27th International Conference on Intelligent User Interfaces. IUI '22 Companion*, pp. 50–53. Association for Computing Machinery, New York, NY, USA (2022). <https://doi.org/>

- [39] Soddu, C.: Generative art geometry. logical interpretations for generative algorithms.
- [40] Farid, H.: Perspective (in) consistency of paint by text. arXiv preprint arXiv:2206.14617 (2022)
- [41] Luccioni, A.S., Akiki, C., Mitchell, M., Jernite, Y.: Stable bias: Analyzing societal representations in diffusion models. arXiv preprint arXiv:2303.11408 (2023)
- [42] Akleman, E., Kurt, M., Akleman, D., Bruins, G., Deng, S., Subramanian, M.: Hyper-realist rendering: A theoretical framework. arXiv preprint arXiv:2401.12853 (2024)
- [43] Foka, A.F.: Computer vision applications for art history: Reflections and paradigms for future research. In: Proceedings of EVA London 2021, pp. 73–80 (2021). BCS Learning & Development
- [44] Castellano, G., Vessio, G.: Deep learning approaches to pattern extraction and recognition in paintings and drawings: An overview. Neural Computing and Applications **33**(19), 12263–12282 (2021)
- [45] Cetinic, E., She, J.: Understanding and creating art with ai: Review and outlook. ACM Transactions on Multimedia Computing, Communications, and Applications (TOMM) **18**(2), 1–22 (2022)
- [46] Remondino, F., Rizzi, A., Barazzetti, L., Scaioni, M., Fassi, F., Brumana, R., Pelagotti, A.: Review of geometric and radiometric analyses of paintings. The Photogrammetric Record **26**(136), 439–461 (2011)
- [47] Castellano, G., Lella, E., Vessio, G.: Visual link retrieval and knowledge discovery in painting datasets. Multimedia Tools and Applications **80**, 6599–6616 (2021)
- [48] Castellano, G., Vessio, G.: Deep convolutional embedding for digitized painting clustering. In: 2020 25th International Conference on Pattern Recognition (ICPR), pp. 2708–2715 (2021). <https://doi.org/10.1109/ICPR48806.2021.9412438>
- [49] Madhu, P., Meyer, A., Zinnen, M., Muhrenberg, L., Suckow, D., Bendschus, T., Reinhardt, C., Bell, P., Verstegen, U., Kosti, R., *et al.*: One-shot object detection in heterogeneous artwork datasets. In: 2022 Eleventh International Conference on Image Processing Theory, Tools and Applications (IPTA), pp. 1–6 (2022). IEEE

- [50] Ahmad, T., Schich, M.: Toward cross-domain object detection in artwork images using improved yolov5 and xgboosting. *IET Image Processing* (2023)
- [51] Fuertes, D., del-Blanco, C.R., Jaureguizar, F., Giarcia, N.: Logomix: A data augmentation technique for object detection applied to logo recognition. In: 2022 IEEE International Conference on Consumer Electronics (ICCE), pp. 1–2 (2022). IEEE
- [52] Smirnov, S., Eguizabal, A.: Deep learning for object detection in fine-art paintings. In: 2018 Metrology for Archaeology and Cultural Heritage (MetroArchaeo), pp. 45–49 (2018). IEEE
- [53] Zhao, Q., Chang, Z., Wang, Z.: Research on the factors affecting accuracy of abstract painting orientation detection. *Multimedia Tools and Applications*, 1–24 (2023)
- [54] Saleh, B., Elgammal, A.: Large-scale classification of fine-art paintings: Learning the right metric on the right feature. *arXiv preprint arXiv:1505.00855* (2015)
- [55] Jeon, H.-J., Jung, S., Choi, Y.-S., Kim, J.W., Kim, J.S.: Object detection in artworks using data augmentation. In: 2020 International Conference on Information and Communication Technology Convergence (ICTC), pp. 1312–1314 (2020). IEEE
- [56] Kadish, D., Risi, S., Lovlie, A.S.: Improving object detection in art images using only style transfer. In: 2021 International Joint Conference on Neural Networks (IJCNN), pp. 1–8 (2021). IEEE
- [57] Sandoval, C., Pirogova, E., Lech, M.: Adversarial learning approach to unsupervised labeling of fine art paintings. *IEEE Access* **9**, 81969–81985 (2021)
- [58] Anwer, R.M., Khan, F.S., Van De Weijer, J., Laaksonen, J.: Combining holistic and part-based deep representations for computational painting categorization. In: Proceedings of the 2016 ACM on International Conference on Multimedia Retrieval, pp. 339–342 (2016)
- [59] Castrejon, L., Aytar, Y., Vondrick, C., Pirsiavash, H., Torralba, A.: Learning aligned cross-modal representations from weakly aligned data. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 2940–2949 (2016)
- [60] Huang, Y., Iizuka, S., Simo-Serra, E., Fukui, K.: Controllable multi-domain semantic artwork synthesis. *arXiv preprint arXiv:2308.10111* (2023)
- [61] Wechsler, H., Toor, A.S.: Modern art challenges face detection. *Pattern Recognition Letters* **126**, 3–10 (2019) <https://doi.org/10.1016/J.PATREC.2018.02.014>

- [62] Yaniv, J.: The face of art: Landmark detection and geometric style in portraits (2019) <https://doi.org/10.1145/3306346.3322984>
- [63] Schneider, S., Vollmer, R.: Poses of people in art: A data set for human pose estimation in digital art history. arXiv preprint arXiv:2301.05124 (2023)
- [64] Ciortan, I.-M., George, S., Hardeberg, J.Y.: Colour-balanced edge-guided digital inpainting: Applications on artworks. *Sensors* **21**(6), 2091 (2021)
- [65] Xue, A.: End-to-end chinese landscape painting creation using generative adversarial networks. In: *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 3863–3871 (2021)
- [66] Bai, Z., Nakashima, Y., Garcia, N.: Explain me the painting: Multi-topic knowledgeable art description generation. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5422–5432 (2021)
- [67] Shen, X., Efros, A.A., Aubry, M.: Discovering visual patterns in art collections with spatially-consistent feature learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 9278–9287 (2019)
- [68] Cai, H., Wu, Q., Corradi, T., Hall, P.: The cross-depiction problem: Computer vision algorithms for recognising objects in artwork and in photographs. arXiv preprint arXiv:1505.00110 (2015)
- [69] Schlecht, J., Carque, B., Ommer, B.: Detecting gestures in medieval image. *IEEE International Conference on Image Processing*, 1285–1288 (2011)
- [70] Gonthier, N., Gousseau, Y., Ladjal, S., Bonfait, O.: Weakly supervised object detection in artworks. In: *Proceedings of the European Conference on Computer Vision (ECCV) Workshops* (2018)
- [71] Lu, Y., Guo, C., Dai, X., Wang, F.Y.: Data-efficient image captioning of fine art paintings via virtual-real semantic alignment training. *Neurocomputing* **490**, 163–180 (2022) <https://doi.org/10.1016/j.neucom.2022.01.068>
- [72] Marinescu, M.-C., Reshetnikov, A., Lopez, J.M.: Improving object detection in paintings based on time contexts. In: *2020 International Conference on Data Mining Workshops (ICDMW)*, pp. 926–932 (2020). IEEE
- [73] Sizyakin, R., Cornelis, B., Meeus, L., Dubois, H., Martens, M., Voronin, V., Pizurica, A.: Crack detection in paintings using convolutional neural networks. *IEEE Access* **8** (2020) <https://doi.org/10.1109/ACCESS.2020.2988856>
- [74] Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, vol. 1, pp. 886–8931 (2005). <https://doi.org/10.1109/CVPR.2005.177>

- [75] Westlake, N., Cai, H., Hall, P.: Detecting people in artwork with cnns. In: Computer Vision–ECCV 2016 Workshops: Amsterdam, The Netherlands, October 8–10 and 15–16, 2016, Proceedings, Part I 14, pp. 825–841 (2016). Springer
- [76] Crowley, E.J., Zisserman, A.: The state of the art: Object retrieval in paintings using discriminative regions. *Proceedings of the British Machine Vision Conference 2014* (2014)
- [77] Lang, S., Ommer, B.: Attesting similarity: Supporting the organization and study of art image collections with computer vision. *Digital Scholarship in the Humanities* **33** (2018) <https://doi.org/10.1093/llc/fqy006>
- [78] Delgado, A., Alba-Carcel'en, L., Murillo-Fuentes, J.J.: Crossing points detection in plain weave for old paintings with deep learning. *arXiv preprint arXiv:2302.11924* (2023)
- [79] Heitzinger, T., Stork, D.G.: Improving semantic segmentation of fine art images using photographs rendered in a style learned from artworks. *Electronic Imaging* **34**(13), 169–11691 (2022) <https://doi.org/10.2352/EL.2022.34.13.CVAA-169>
- [80] Zhang, J., Liu, C., Xian, K., Cao, Z.: Large motion anime head animation using a cascade pose transform network. *Pattern Recognition* **135** (2023) <https://doi.org/10.1016/J.PATCOG.2022.109181>
- [81] Springstein, M., Schneider, S., Althaus, C., Ewerth, R., Ew, R.: Semi-supervised human pose estimation in art-historical images. In: *Proceedings of the 30th ACM International Conference on Multimedia (MM '22)*, Oct, 2022, Lisboa, Portugal, vol. 1 (2022). <https://doi.org/10.1145/3503161.3548371> . <https://doi.org/10.1145/3503161.3548371>
- [82] Sindel, A., Maier, A., Christlein, V.: Artfacepoints: High-resolution facial landmark detection in paintings and prints (2022)
- [83] Carneiro, G., Da Silva, N.P., Del Bue, A., Costeira, J.P.: Artistic image classification: An analysis on the printart database. In: *Computer Vision–ECCV 2012: 12th European Conference on Computer Vision, Florence, Italy, October 7–13, 2012, Proceedings, Part IV* 12, pp. 143–157 (2012). Springer
- [84] Ju, X., Zeng, A., Wang, J., Xu, Q., Zhang, L.: Human-art: A versatile human-centric dataset bridging natural and artificial scenes. *Computer Vision and Pattern Recognition* (2023) <https://doi.org/10.1109/CVPR52729.2023.00067>
- [85] Nguyen, D.T., Li, W., Ogunbona, P.O.: Human detection from images and videos: A survey. *Pattern Recognition* **51**, 148–175 (2016)
- [86] Hall, P., Cai, H., Wu, Q., Corradi, T.: Cross-depiction problem: Recognition and synthesis of photographs and artwork. *Computational Visual Media* **1**, 91–103

(2015)

- [87] Arkin, E., Yadikar, N., Xu, X., Aysa, A., Ubul, K., Tools, M.: A survey: object detection methods from cnn to transformer. *Multimedia Tools and Applications* (2023) <https://doi.org/10.1007/s11042-022-13801-3>
- [88] Wang, X., Girdhar, R., Yu, S.X., Misra, I.: Cut and learn for unsupervised object detection and instance segmentation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3124–3134 (2023)
- [89] Liu, Y., Yang, D., Wang, Y., Liu, J., Liu, J., Boukerche, A., Sun, P., Song, L.: Generalized video anomaly event detection: Systematic taxonomy and comparison of deep models. *ACM Computing Surveys* (2023)
- [90] Liu, Y., Liu, J., Yang, K., Ju, B., Liu, S., Wang, Y., Yang, D., Sun, P., Song, L.: Amp-net: Appearance-motion prototype network assisted automatic video anomaly detection system. *IEEE Transactions on Industrial Informatics* (2023)
- [91] Liu, Y., Xia, Z., Zhao, M., Wei, D., Wang, Y., Liu, S., Ju, B., Fang, G., Liu, J., Song, L.: Learning causality-inspired representation consistency for video anomaly detection. In: *Proceedings of the 31st ACM International Conference on Multimedia*, pp. 203–212 (2023)
- [92] Liu, J., Liu, Y., Zhu, W., Zhu, X., Song, L.: Distributional and spatial-temporal robust representation learning for transportation activity recognition. *Pattern Recognition* **140**, 109568 (2023)
- [93] Vijendran, M., Li, F.W.B., Shum, H.P.H.: Tackling data bias in painting classification with style transfer. In: *Proceedings of the 2023 International Conference on Computer Vision Theory and Applications. VISAPP '23*, pp. 250–261 (2023). <https://doi.org/10.5220/0011776600003417>
- [94] Feng, Y., Jiang, J., Tang, M., Jin, R., Gao, Y.: Rethinking supervised pre-training for better downstream transferring. *arXiv preprint arXiv:2110.06014* (2021)
- [95] Smirnov, S., Eguizabal, A.: Deep learning for object detection in fine-art paintings. In: *2018 Metrology for Archaeology and Cultural Heritage (MetroArchaeo)*, pp. 45–49 (2018). IEEE
- [96] Li, D., Yang, Y., Song, Y.-Z., Hospedales, T.M.: Deeper, broader and artier domain generalization. In: *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5542–5550 (2017)
- [97] Lazzeri, D., Nicoli, F., Zhang, Y.X.: Secret hand gestures in paintings. *Acta Bio Medica: Atenei Parmensis* **90**(4), 526 (2019)

- [98] Zheng, C., Wu, W., Chen, C., Yang, T., Zhu, S., Shen, J., Kehtarnavaz, N., Shah, M.: Deep learning-based human pose estimation: A survey. *ACM Computing Surveys* **56**(1), 1–37 (2023)
- [99] Wang, J., Sun, K., Cheng, T., Jiang, B., Deng, C., Zhao, Y., Liu, D., Mu, Y., Tan, M., Wang, X., Liu, W., Xiao, B.: Deep high-resolution representation learning for visual recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **43**(10), 3349–3364 (2021) <https://doi.org/10.1109/TPAMI.2020.2983686>
- [100] He, N., Lu, K.: An image segmentation method for chinese paintings by combining deformable models with graph cuts. In: *Human-Computer Interaction. Design and Development Approaches: 14th International Conference, HCI International 2011, Orlando, FL, USA, July 9-14, 2011, Proceedings, Part I* 14, pp. 571–579 (2011). Springer
- [101] Kamann, C., Rother, C.: Increasing the robustness of semantic segmentation models with painting-by-numbers. In: *European Conference on Computer Vision*, pp. 369–387 (2020). Springer
- [102] Chen, Y., Yuan, Q., Li, Z., Xie, Y.L.W.W.C., Wen, X., Yu, Q.: Upst-nerf: Universal photorealistic style transfer of neural radiance fields for 3d scene. *arXiv preprint arXiv:2208.07059* (2022)
- [103] Carroll, R., Agarwala, A., Agrawala, M.: Image warps for artistic perspective manipulation. In: *ACM SIGGRAPH 2010 Papers*, pp. 1–9 (2010)
- [104] Sahay, P., Rajagopalan, A.: Geometric inpainting of 3d structures. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pp. 1–7 (2015)
- [105] Kim, Y., Winnemoller, H., Lee, S.: Wysiwyg stereo painting. In: *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*, pp. 169–176 (2013)
- [106] Huang, Y.-H., He, Y., Yuan, Y.-J., Lai, Y.-K., Gao, L.: Stylizednerf: consistent 3d scene stylization as stylized nerf via 2d-3d mutual learning. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18342–18352 (2022)
- [107] Tseng, K.-W., Lee, Y.-C., Chen, C.-S.: Artistic style novel view synthesis based on a single image. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2258–2262 (2022)
- [108] Zhang, K., Kolkin, N., Bi, S., Luan, F., Xu, Z., Shechtman, E., Snavely, N.: Arf: Artistic radiance fields. In: *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXI*, pp.

- [109] Srinivasan, P.P., Deng, B., Zhang, X., Tancik, M., Mildenhall, B., Barron, J.T.: Nerv: Neural reflectance and visibility fields for relighting and view synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 7495–7504 (2021)
- [110] Chan, E.R., Lin, C.Z., Chan, M.A., Nagano, K., Pan, B., De Mello, S., Gallo, O., Guibas, L.J., Tremblay, J., Khamis, S., *et al.*: Efficient geometry-aware 3d generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 16123–16133 (2022)
- [111] Chang, Z., Koulrieris, G.A., Shum, H.P.: 3d reconstruction of sculptures from single images via unsupervised domain adaptation on implicit models. In: Proceedings of the 28th ACM Symposium on Virtual Reality Software and Technology, pp. 1–10 (2022)
- [112] Pang, S., Peng, R., Dong, Y., Yuan, Q., Wang, S., Sun, J.: Jointmetro: a 3d reconstruction model for human figures in works of art based on transformer. *Neural Computing and Applications*, 1–15 (2023)
- [113] Casati, P., Ronfard, R., Hahmann, S.: Approximate reconstruction of 3d scenes from bas-reliefs. In: GCH 2019-EUROGRAPHICS Workshop on Graphics and Cultural Heritage, pp. 109–118 (2019). The Eurographics Association
- [114] Zeidler, D., McGinity, M.: Bodylab: in virtuo sculpting, painting and performing of full-body avatars. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* **6**(2), 1–12 (2023)
- [115] Jetchev, N.: Clipmatrix: Text-controlled creation of 3d textured meshes. *arXiv preprint arXiv:2109.12922* (2021)
- [116] Fu, T., Chaine, R., Digne, J.: Fakir: An algorithm for revealing the anatomy and pose of statues from raw point sets. In: *Computer Graphics Forum*, vol. 39, pp. 375–385 (2020). Wiley Online Library
- [117] Chang, Z., Koulrieris, G.A., Shum, H.P.H.: 3d reconstruction of sculptures from single images via unsupervised domain adaptation on implicit models (2022)
- [118] Thomas, C., Kovashka, A.: Artistic object recognition by unsupervised style adaptation. In: *Computer Vision–ACCV 2018: 14th Asian Conference on Computer Vision*, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part III 14, pp. 460–476 (2019). Springer
- [119] Seo, S., Lee, H., Kim, Y., Son, W.: Video motion analysis for landscape image abstraction. In: *2016 International Conference on Platform Technology and Service (PlatCon)*, pp. 1–4 (2016). IEEE

- [120] Pasqualino, G., Furnari, A., Farinella, G.M.: A multi camera unsupervised domain adaptation pipeline for object detection in cultural sites through adversarial learning and self-training. *Computer Vision and Image Understanding* **222**, 103487 (2022)
- [121] Condorovici, R.G., Florea, C., Vertan, C.: Painting scene recognition using homogenous shapes. In: *Advanced Concepts for Intelligent Vision Systems: 15th International Conference, ACIVS 2013, Poznan, Poland, October 28-31, 2013. Proceedings 15*, pp. 262–273 (2013). Springer
- [122] Todorovic, D.: The effect of the observer vantage point on perceived distortions in linear perspective images. *Attention, Perception, and Psychophysics* **71**, 183–193 (2009) <https://doi.org/10.3758/APP.71.1.183>
- [123] Rapp, J.B.: A geometrical analysis of multiple viewpoint perspective in the work of giovanni battista piranesi: an application of geometric restitution of perspective. *The Journal of architecture* **13**(6), 701–736 (2008)
- [124] Fumanal-Idocin, J., Andreu-Perez, J., Cordon, O., Hagraas, H., Bustince, H.: Artxai: Explainable artificial intelligence curates deep representation learning for artistic images using fuzzy techniques. *arXiv preprint arXiv:2308.15284* (2023)
- [125] Liu, X.-C., Yang, Y.-L., Hall, P.: Geometric and textural augmentation for domain gap reduction. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14340–14350 (2022)
- [126] Yang, J., Guo, F., Chen, S., Li, J., Yang, J.: Industrial style transfer with large-scale geometric warping and content preservation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7834–7843 (2022)
- [127] Liu, X.-C., Yang, Y.-L., Hall, P.: Learning to warp for style transfer. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3702–3711 (2021)
- [128] Vulimiri, P.S., Deng, H., Dugast, F., Zhang, X., To, A.C.: Integrating geometric data into topology optimization via neural style transfer. *Materials* **14**(16), 4551 (2021)
- [129] Nakano, R.: Neural painters: A learned differentiable constraint for generating brushstroke paintings. *arXiv preprint arXiv:1904.08410* (2019)
- [130] Geng, J., Ma, L., Li, X., Yan, Y.: Ptgcfc: Printing texture guided color fusion for impressionism oil painting style rendering. *arXiv e-prints*, 2207 (2022)
- [131] Papari, G., Petkov, N.: Glass patterns and artistic imaging. In: *Advances in Image and Video Technology: Third Pacific Rim Symposium, PSIVT 2009*,

- Tokyo, Japan, January 13-16, 2009. Proceedings 3, pp. 1034–1045 (2009). Springer
- [132] Yin, W., Liu, Z., Loy, C.C.: Instance-level facial attributes transfer with geometry-aware flow. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 33, pp. 9111–9118 (2019)
 - [133] Kim, S.S., Kolkin, N., Salavon, J., Shakhnarovich, G.: Deformable style transfer. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXVI 16, pp. 246–261 (2020). Springer
 - [134] Kopanas, G., Philip, J., Leimkuhler, T., Drettakis, G.: Point-based neural rendering with per-view optimization. In: Computer Graphics Forum, vol. 40, pp. 29–43 (2021). Wiley Online Library
 - [135] Liu, X.-C., Li, X.-Y., Cheng, M.-M., Hall, P.: Geometric style transfer. arXiv preprint arXiv:2007.05471 (2020)
 - [136] Alexandru, I., Nicula, C., Prodan, C., Rotaru, R.-P., Voncilua, M.-L., Tarbua, N., Boiangiu, C.-A.: Image style transfer via multi-style geometry warping. *Applied Sciences* **12**(12), 6055 (2022)
 - [137] Du, X., He, Y., Yang, X., Chang, C.-M., Xie, H.: Sketch-based 3d shape modeling from sparse point clouds. arXiv, 119 (2022) <https://doi.org/10.1117/12.2626116>
 - [138] Upadhyay, A., Dubey, A., Kuriakose, S.M., Mahato, D.: 3dstnet: Neural 3d shape style transfer. In: 2022 IEEE International Conference on Multimedia and Expo Workshops (ICMEW), pp. 1–6 (2022). IEEE
 - [139] Bird, J.J.: Continuation of famous art with ai: A conditional adversarial network inpainting approach. arXiv preprint arXiv:2110.09170 (2021)
 - [140] Cipolina-Kun, L., Papadakis, S.M., Caenazzo, S.: Discriminative candidate selection for image inpainting applications to the fine arts. *LatinX in AI at International Conference on Machine Learning 2022* (2022) <https://doi.org/10.52591/lxai202207176>
 - [141] Zhang, L., Agrawala, M.: Adding Conditional Control to Text-to-Image Diffusion Models (2023)
 - [142] Chen, X., Jin, X., Zhao, Q., Wu, H.: Artistic illumination transfer for portraits. In: Computer Graphics Forum, vol. 31, pp. 1425–1434 (2012). Wiley Online Library
 - [143] Chen, W.-Y., Ople, J.J.M., Si, M.J., Tan, D.S., Hua, K.-L.: Perspective preserving style transfer for interior portraits. *IEEE Access* **9**, 7033–7042 (2021)

- [144] Henz, B., Oliveira, M.M.: Artistic relighting of paintings and drawings. *The Visual Computer* **33**(1), 33–46 (2017) <https://doi.org/10.1007/s00371-015-1150-7>
- [145] Henz, B.: Image relighting using shading proxies. (2014)
- [146] Mishra, S., Granskog, J.: Clip-based neural neighbor style transfer for 3d assets. *ArXiv abs/2208.04370* (2022)
- [147] Jin, B., Tian, B., Zhao, H., Zhou, G.: Language-guided semantic style transfer of 3d indoor scenes. *Proceedings of the 1st Workshop on Photorealistic Image and Environment Synthesis for Multimedia Experiments* (2022)
- [148] Zhao, A., Balakrishnan, G., Lewis, K.M., Durand, F., Gutttag, J.V., Dalca, A.V.: Painting many pasts: Synthesizing time lapse videos of paintings. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 8435–8445 (2020)
- [149] Tan, J., Dvorožňák, M., Šykora, D., Gingold, Y.: Decomposing time-lapse paintings into layers. *ACM Transactions on Graphics (TOG)* **34**(4), 1–10 (2015)
- [150] Koyama, Y., Goto, M.: Decomposing images into layers with advanced color blending. In: *Computer Graphics Forum*, vol. 37, pp. 397–407 (2018). Wiley Online Library
- [151] Hou, H., Huo, J., Wu, J., Lai, Y.-K., Gao, Y.: Mw-gan: multi-warping gan for caricature generation with multi-style geometric exaggeration. *IEEE Transactions on Image Processing* **30**, 8644–8657 (2021)
- [152] Abrahamsen, N., Yao, J.: Inventing painting styles through natural inspiration. *arXiv preprint arXiv:2305.12015* (2023)
- [153] Chang, Z., Kouliris, G.A., Shum, H.P.H.: On the design fundamentals of diffusion models: A survey. *arXiv preprint arXiv: 2306.04542* (2023)
- [154] Chen, D.-Y.: Conditional human sketch synthesis with explicit abstraction control. *arXiv preprint arXiv:2306.09274* (2023)
- [155] Peng, Y., Zhao, C., Xie, H., Fukusato, T., Miyata, K.: Difffacesketch: High-fidelity face image synthesis with sketch-guided latent diffusion model. *arXiv preprint arXiv:2302.06908* (2023)
- [156] Zeng, Y., Lin, Z., Zhang, J., Liu, Q., Collomosse, J., Kuen, J., Patel, V.M.: Scenecomposer: Any-level semantic image synthesis. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22468–22478 (2023)
- [157] Kolkin, N., Salavon, J., Shakhnarovich, G.: Style transfer by relaxed optimal

- transport and self-similarity. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 10051–10060 (2019)
- [158] Datta, R., Ghorai, M., Mandal, S.: Image inpainting using geometric transformations for digital circuit images. In: 2017 Ninth International Conference on Advances in Pattern Recognition (ICAPR), pp. 1–6 (2017). IEEE
 - [159] Ci, Y., Ma, X., Wang, Z., Li, H., Luo, Z.: User-guided deep anime line art colorization with conditional adversarial networks. In: Proceedings of the 26th ACM International Conference on Multimedia, pp. 1536–1544 (2018)
 - [160] He, B., Gao, F., Ma, D., Shi, B., Duan, L.-Y.: Chipgan: A generative adversarial network for chinese ink wash painting style transfer. In: Proceedings of the 26th ACM International Conference on Multimedia, pp. 1172–1180 (2018)
 - [161] Chen, M., Laina, I., Vedaldi, A.: Training-free layout control with cross-attention guidance. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision, pp. 5343–5353 (2024)
 - [162] Abate, D.: Documentation of paintings restoration through photogrammetry and change detection algorithms. *Heritage Science* **7** (2019) <https://doi.org/10.1186/s40494-019-0257-y>
 - [163] Castagnetti, C., Rossi, P., Capra, A.: 3d reconstruction of rock paintings: A cost-effective approach based on modern photogrammetry for rapidly mapping archaeological findings, vol. 364 (2018). <https://doi.org/10.1088/1757-899X/364/1/012020>
 - [164] Carrozzino, M., Evangelista, C., Brondi, R., Tecchia, F., Bergamasco, M.: Virtual reconstruction of paintings as a tool for research and learning. *Journal of Cultural Heritage* **15**, 308–312 (2014) <https://doi.org/10.1016/j.culher.2013.06.003>
 - [165] Bent, G.R., Pfaff, D., Brooks, M., Radpour, R., Delaney, J.: A practical workflow for the 3d reconstruction of complex historic sites and their decorative interiors: Florence as it was and the church of orsanmichele. *Heritage Science* **10** (2022) <https://doi.org/10.1186/s40494-022-00750-1>
 - [166] Jackson, A.S., Bulat, A., Argyriou, V., Tzimiropoulos, G.: Large pose 3d face reconstruction from a single image via direct volumetric cnn regression. In: Proceedings of the IEEE International Conference on Computer Vision, vol. 2017-October, pp. 1031–1039 (2017). <https://doi.org/10.1109/ICCV.2017.117>
 - [167] Zhou, X., In, D., Chen, X., Liu, X., Yang, Y.: Spectral 3d reconstruction of impressionist oil painting based on macroscopic oct imaging. *Optics InfoBase Conference Papers* (2020) <https://doi.org/10.1364/ao.390326>

- [168] Moradi, M., Ghorbani, R., Sfarra, S., Tax, D.M.J., Zarouchas, D.: A spatiotemporal deep neural network useful for defect identification and reconstruction of artworks using infrared thermography. *Sensors* **22** (2022) <https://doi.org/10.3390/s22239361>
- [169] Léang, M., Giorgiutti-Dauphiné, F., Lee, L.T., Pauchard, L.: Crack opening: From colloidal systems to paintings. *Soft Matter* **13**, 5802–5808 (2017) <https://doi.org/10.1039/c7sm00985b>
- [170] Yuan, J., Chen, C., Yao, D., Chen, G.: 3d printing of oil paintings based on material jetting and its reduction of staircase effect. *Polymers* **12**, 1–12 (2020) <https://doi.org/10.3390/polym12112536>
- [171] Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III* 18, pp. 234–241 (2015). Springer
- [172] Barazzetti, L., Remondino, F., Scaioni, M., Lo Brutto, M., Rizzi, A., Brumana, R., et al.: Geometric and radiometric analysis of paintings. *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences* **38**(Part 5) (2010)
- [173] Blais, F., Taylor, J., Cournoyer, L., Picard, M., Borgeat, L., Dicaire, L., Rioux, M., Beraldin, J., Godin, G., Lahanier, C., et al.: Ultra-high resolution imaging at 50 μ m using a portable xyz-rgb color laser scanner. In: *International Workshop on Recording, Modeling and Visualization of Cultural Heritage*, p. 48099 (2005). NRC Ascona, Switzerland
- [174] Lee, J., Kim, E., Lee, Y., Kim, D., Chang, J., Choo, J.: Reference-based sketch image colorization using augmented-self reference and dense semantic correspondence. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 5800–5809 (2020)
- [175] Schaldenbrand, P., Oh, J.: Content masked loss: Human-like brush stroke planning in a reinforcement learning painting agent. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, pp. 505–512 (2021)
- [176] Zhang, Y., Zhang, Z., DiVerdi, S., Wang, Z., Echevarria, J., Fu, Y.: Texture hallucination for large-factor painting super-resolution. In: *European Conference on Computer Vision*, pp. 209–225 (2020). Springer
- [177] Yuan, S., Dai, A., Yan, Z., Liu, R., Chen, M., Chen, B., Qiu, Z., He, X.: Learning to generate poetic chinese landscape painting with calligraphy. *arXiv preprint arXiv:2305.04719* (2023)

- [178] Shahid, M., Koch, M., Schneider, N.: Paint it black: Generating paintings from text descriptions. arXiv preprint arXiv:2302.08808 (2023)
- [179] Tong, Z., Wang, X., Yuan, S., Chen, X., Wang, J., Fang, X.: Im2oil: Stroke-based oil painting rendering with linearly controllable fineness via adaptive sampling. In: Proceedings of the 30th ACM International Conference on Multimedia, pp. 1035–1046 (2022)
- [180] Huang, Y., Iizuka, S., Simo-Serra, E., Fukui, K.: Controllable multi-domain semantic artwork synthesis. *Computational Visual Media* **10**(2), 355–373 (2024)
- [181] Singh, J., Zheng, L., Smith, C., Echevarria, J.: Paint2pix: interactive painting based progressive image synthesis and editing. In: European Conference on Computer Vision, pp. 678–695 (2022). Springer
- [182] Lourakis, M., Alongi, P., Delouis, D., Lippi, F., Spadoni, F., SpA, P.A.S.: RECOVER: PHOTOREALISTIC 3D RECONSTRUCTION OF PERSPECTIVE PAINTINGS AND PICTURES. <http://www.ics.forth.gr/recover/>.
- [183] Li, X., Lin, C.-C., Chen, Y., Liu, Z., Wang, J., Raj, B.: Paintseg: Training-free segmentation via painting. arXiv preprint arXiv:2305.19406 (2023)
- [184] Singh, J., Zheng, L.: Combining semantic guidance and deep reinforcement learning for generating human level paintings. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 16387–16396 (2021)
- [185] Han, X., Wu, Y., Wan, R.: A method for style transfer from artistic images based on depth extraction generative adversarial network. *Applied Sciences* **13**(2), 867 (2023)
- [186] Yang, B., Zhang, Y., Xu, Y., Li, Y., Zhou, H., Bao, H., Zhang, G., Cui, Z.: Learning object-compositional neural radiance field for editable scene rendering. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 13779–13788 (2021)
- [187] Park, K., Sinha, U., Barron, J.T., Bouaziz, S., Goldman, D.B., Seitz, S.M., Martin-Brualla, R.: Nerfies: Deformable neural radiance fields. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 5865–5874 (2021)
- [188] Zhao, Y., Barnes, C., Zhou, Y., Shechtman, E., Amirghodsi, S., Fowlkes, C.: Geofill: Reference-based image inpainting with better geometric understanding. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), pp. 1776–1786 (2023)
- [189] Sofiiuk, K., Petrov, I.A., Konushin, A.: Reviving iterative training with mask

- guidance for interactive segmentation. In: 2022 IEEE International Conference on Image Processing (ICIP), pp. 3141–3145 (2022). IEEE
- [190] Bragantini, J., Falcão, A.X., Najman, L.: Rethinking interactive image segmentation: Feature space annotation. *Pattern Recognition* **131**, 108882 (2022)
 - [191] Groenen, I., Rudinac, S., Worring, M.: Panorams: automatic annotation for detecting objects in urban context. *IEEE Transactions on Multimedia* (2023)
 - [192] Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.-Y., *et al.*: Segment anything. In: Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 4015–4026 (2023)
 - [193] Wang, X., Guo, P., Zhang, Y.: Domain adaptation via bidirectional cross-attention transformer. *arXiv preprint arXiv:2201.05887* (2022)
 - [194] Zhang, X., Chen, Y., Shen, Z., Shen, Y., Zhang, H., Zhang, Y.: Confidence-and-refinement adaptation model for cross-domain semantic segmentation. *IEEE Transactions on Intelligent Transportation Systems* **23**(7), 9529–9542 (2022)
 - [195] Hertz, A., Mokady, R., Tenenbaum, J., Aberman, K., Pritch, Y., Cohen-Or, D.: Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626* (2022)
 - [196] Zhao, S., Chen, D., Chen, Y.-C., Bao, J., Hao, S., Yuan, L., Wong, K.-Y.K.: Uni-controlnet: All-in-one control to text-to-image diffusion models. *arXiv preprint arXiv:2305.16322* (2023)
 - [197] Wu, C.H., Torre, F.: Unifying diffusion models’ latent space, with applications to cyclediffusion and guidance. *arXiv preprint arXiv:2210.05559* (2022)
 - [198] Tertikas, K., Paschalidou, D., Pan, B., Park, J.J., Uy, M.A., Emiris, I., Avrithis, Y., Guibas, L.: Generating part-aware editable 3d shapes without 3d supervision. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 4466–4478 (2023)
 - [199] Aygün, M., Mac Aodha, O.: Saor: Single-view articulated object reconstruction. *arXiv preprint arXiv:2303.13514* (2023)
 - [200] Zheng, X.-Y., Pan, H., Wang, P.-S., Tong, X., Liu, Y., Shum, H.-Y.: Locally attentional sdf diffusion for controllable 3d shape generation. *arXiv preprint arXiv:2305.04461* (2023)