








AID-AppEAL: Automatic Image Dataset and Algorithm for Content Appeal Enhancement and Assessment Labeling

Sherry X. Chen^{1*}, Yaron Vaxman², Elad Ben Baruch², David Asulin²,
Aviad Moreshet², Misha Sra¹, and Pradeep Sen¹

¹ University of California, Santa Barbara

² Clouidary

Abstract. We propose Image Content Appeal Assessment (ICAA), a novel metric that quantifies the level of positive interest an image’s content generates for viewers, such as the appeal of food in a photograph. This is fundamentally different from traditional Image-Aesthetics Assessment (IAA), which judges an image’s artistic quality. While previous studies often confuse the concepts of “aesthetics” and “appeal,” our work addresses this by being the first to study ICAA explicitly. To do this, we propose a novel system that automates dataset creation and implements algorithms to estimate and boost content appeal. We use our pipeline to generate two large-scale datasets (70K+ images each) in diverse domains (food and room interior design) to train our models, which revealed little correlation between content appeal and aesthetics. Our user study, with more than 76% of participants preferring the appeal-enhanced images, confirms that our appeal ratings accurately reflect user preferences, establishing ICAA as a unique evaluative criterion. Our code and datasets are available at <https://github.com/SherryXTChen/AID-Appeal>.

Keywords: image assessment · automated dataset creation · image manipulation

1 Introduction

The accurate measure of perceptual image quality is an important problem in computer vision, since algorithms must often account for how humans actually perceive images. For this reason, researchers have developed algorithms focusing on distinct facets of image quality. For example, image-quality assessment (IQA) algorithms aim to estimate the perceptual impact of distortions [12, 17, 19, 21, 24, 26, 37, 42, 60, 62], while image aesthetics assessment (IAA) evaluates an image’s aesthetics based on principles of art and photography [18, 30, 33, 46, 53, 57, 63, 65].

Beyond IQA and IAA, we identify a critical yet overlooked aspect of perceptual image quality: image-content appeal assessment (ICAA). This concept

* Corresponding author email: xchen774@ucsb.edu



DIAA	<u>0.63</u>	0.65	0.68	<u>0.34</u>	0.42	0.36
MPADA	<u>0.44</u>	0.61	0.49	<u>0.48</u>	0.70	0.74
NIMA	<u>5.51</u>	6.13	6.27	<u>5.68</u>	5.95	5.98
Ours	7.80	<u>5.52</u>	8.08 (+2.56)	7.70	<u>3.51</u>	6.64 (+3.13)

Fig. 1: Image-content appeal assessment (ICAA) and enhancement. The 1st/4th columns show amateur photos lacking artistic appeal, while the 2nd/5th columns feature professionally taken images of less appealing content (a moldy burger and a dirty room). Because of their superior aesthetics, IAA baselines (DIAA [22], MPADA [53], and NIMA [57]) rate them higher even though they have less appealing content (lowest scores underlined, highest in bold), while our ICAA estimator accurately assesses and enhances content appeal (2nd/5th to 3rd/6th columns).

becomes evident when comparing professional photographs that are highly aesthetic (Fig. 1, Cols.2,5) but feature unappealing subjects (a moldy burger and a dirty room). As these images receive high scores from existing IAA methods, including some designed to measure “image appeal,” it underscores the need to evaluate a quality fundamentally different from existing metrics. We call it “image content appeal.”

To devise a formal definition of image-content appeal, we take inspiration from the photography literature, where image appeal is defined as “the interest that a picture generates when viewed by third-party observers” [51]. Our focus, however, shifts from the image itself to the content it portrays, emphasizing the amount of **positive interest in the content of a picture when viewed by generic third-party observers**. This distinction allows us to assess how much a viewer might desire to engage with the content, such as eating the food shown in the picture or staying in a depicted room. A metric like this would benefit sectors like food services, online retail, and vacation rentals, for example.

Image-content appeal assessment (ICAA) emerges as a compelling research avenue with significant practical implications. However, the absence of dedicated datasets for ICAA research presents a challenge, as existing image aesthetics assessment (IAA) datasets only broadly cover “interesting content” [18, 22, 23, 65] or “interest-ness” [10, 30, 61], not specifically targeting the positive interest ICAA focuses on. Another option is to create our own ICAA dataset, but manually annotating large image assessment datasets (IQA [12, 14, 19, 20, 32, 40, 41, 52, 62]; IAA [8, 9, 16, 36, 46, 61]) can become an expensive and time-consuming bottleneck.

To bridge this gap, we present AID-AppEAL, an automated dataset generation pipeline as well as algorithms for estimating and enhancing content appeal. We used our system to generate two large-scale datasets (food and room interior design), totaling over 70,000 images each, which enable the training of specialized content appeal estimators and enhancers. Our content appeal scores show little correlation with traditional aesthetics scores, underscoring the distinct nature of

ICAA. User studies further validate our approach, with over 76% of participants favoring the appeal-enhanced images, affirming the effectiveness of our system in accurately capturing and enhancing image-content appeal.

In summary, the main contributions of our work are:

1. Recognition of image-content appeal (ICAA) as distinct from traditional image-aesthetic and appeal assessments that have been previously studied
2. Development of a universal automated ICAA dataset creation pipeline
3. Creation of two domain-specific ICAA datasets
4. Introduction of accurate ICAA estimators for each of the two datasets
5. Implementation of content appeal enhancers for each dataset, improving ICAA while maintaining visual integrity, validated by a user study.

2 Related Work

2.1 Image aesthetics and content appeal assessment

Previous research in image aesthetics and appeal, collectively termed image aesthetic assessment (IAA), aims to evaluate an image’s quality and attractiveness. In those works, the term “aesthetic appeal” has been used to denote “the subjective notions of “beauty” in the image [45], “what makes an image aesthetically pleasing” [55], and “appeal” is the quality of the image “being attractive or interesting” [15]. ‘Appeal’ in these contexts refers to how appealing images are from an artistic point of view.

Since different factors such as lighting, contrast, color harmony, and composition all play a role in assessing image aesthetics, prior IAA methods often design different branches in their neural networks that either take different crops of each image [30, 33] or estimate various IAA attribute scores [46, 53, 63]. Convolution neural networks (CNNs) followed by fully connected layers (FCs) are commonly architectural components used as the backbone of these algorithms [18, 30, 33, 46, 53, 57, 63, 65]. Some work also adapts variations of pre-trained visual models to facilitate their tasks [57, 63, 65].

2.2 Image assessment dataset creation

Creating a good dataset is often one of the most critical steps for image assessment research. In image-quality assessment (IQA), the goal is to evaluate the perceived technical quality of an image after it is distorted. IQA datasets can either be full-reference (FR-IQA) or no-reference (NR-IQA), depending on whether the original pristine reference images are available in the dataset.

FR-IQA datasets are created from a set of pristine images, where various distortion operations are applied to create different distorted versions of them [20, 28, 32, 40–42, 52]. While these datasets are mostly human-annotated, this dataset creation process does offer some leeway to annotate images automatically based on the distortion operations being applied. On the other hand,

samples in these datasets are heavily correlated. Furthermore, the distortion may not fully reflect the characteristics of distorted image “in the wild.”

In contrast, NR-IQA datasets contain distorted images where their pristine counterparts are unknown, often because we only have the final distorted images from the internet [12, 14, 19, 62]. As a result, these datasets are usually much larger, more diverse, and contain more realistic samples. On the other hand, extensive laboratory subjective study [12] or crowdsourcing [14, 19, 62] is required for dataset annotation, which is expensive and time-consuming.

Another type of image assessment is image aesthetics assessment (IAA), which is concerned with human perception of beauty. Although this is largely a subjective manner, some prior work [8, 9, 22, 36] annotated images by having multiple annotators labeling one image so “the average score can be thought of as an estimator for its intrinsic aesthetic quality” [9]. On the other hand, more recent work [16, 46, 61] acknowledge the subjective nature of aesthetics assessment and focus on not just the average score, but scores from each annotator to study and estimate human personal preference.

2.3 Image generation and appeal enhancement

Both our ICAA dataset creation pipeline and our content appeal enhancement module are made possible due to recent advances in image generation, in particular with GAN-based models [29, 48–50], diffusion-based models [3, 5, 11, 38, 39, 47], and even some combinations of them in between [59]. There is also an increasing interest in using these models to create training datasets [2, 5, 7, 54].

Although there isn’t any work dedicated specifically to content appeal enhancement to the best of our knowledge, there are a lot of image manipulation methods that have the potential to enable such applications, where including text-based image editing [3, 5, 39], local inpainting with additional user-defined mask [1, 34, 44], invert and finetuning pre-trained models on an image-text pair and edit the text to enable localized editing [11, 35, 56]. Our content appeal enhancement method has borrowed components from these methods and opted to use an automated generated mask and textual inversion [13] for image editing.

3 Dataset Creation Pipeline

To develop effective ICAA algorithms, it’s crucial to assemble a dataset that meets specific criteria: it should align with human perception of content appeal, span a broad spectrum of appeal levels and content variety to avoid overfitting, and include a significant volume of high-resolution images for robust machine learning model training. Our initial research indicated that consumer photos typically exhibit a limited range of appeal, often skewing towards the lower end. To counteract potential bias, we exclusively incorporated professional images to ensure both high aesthetic and quality standards, enabling our model to accurately assess content appeal. The trained models can also generalize to consumer-taken pictures, as shown in Sec. 5.




Fig. 2: Domain-relevancy map generation. Given an image, we use BLIP [27] to estimate its description and extract all noun phrases \mathbb{P} using NLTK [4]. For every phrase, we look up each of its words in WordNet [6] to get their lexnames and keep the phrase if any of them matches the domain D (e.g., if D is food, then the phrase is kept only if at least one word’s lexname is *noun.food*). The resulting set of phrases is \mathbb{P}_D and we use CLIPSeg [31] to create a segmentation map that locates objects described by each phrase in \mathbb{P}_D . These maps collectively define the image region that contains objects from D , and we call it the domain-relevancy map.

Creating such a dataset poses substantial challenges, including the costly manual labeling process, the difficulty in accessing a large volume of high-quality professional images due to stock photo website restrictions, and the inherent bias towards appealing content on these platforms. For instance, a search for “delicious burger” yields hundreds of thousands of results on major stock image sites, whereas “moldy burger” returns vastly fewer, risking biased model training.

To address these issues, we introduce an automated pipeline for generating domain-specific ICAA datasets, which ensure domain consistency to maintain relevance (e.g., food, rooms, scenery). Our approach involves collecting a base set of domain-matching images from stock websites, processing them to highlight domain-specific elements, and creating a synthetic dataset through image manipulation to vary content appeal and background diversity. This dataset trains a relative content appeal comparator to label a vast collection of real images for the final dataset, facilitating the efficient creation of large-scale datasets (over 70K images per domain) without manual labeling. While our examples focus on food imagery, this methodology is adaptable to various image domains, as further detailed in the supplementary.

3.1 Base image set collection and pre-processing

To construct a comprehensive ICAA dataset tailored to a specific application domain D (e.g., food, room interiors, scenery, people), we start by gathering a modest collection of domain-specific real images to generate a synthetic dataset. This synthetic dataset is then utilized to train an automatic labeling system, which produces the final dataset.

To automatically generate search queries for stock-image websites to retrieve suitable images, the process begins with defining a set of nouns \mathbb{N}_D that represent elements within the domain D , such as {“burger,” “cake,” “fruit,” ...} for food.

We then identify two sets of adjectives, \mathbb{A}_D^+ or positive (appealing) descriptors (e.g., {“delicious,” “gourmet,” “tasty,” ...} for food) and \mathbb{A}_D^- for negative (unappealing) descriptors (e.g., {“disgusting,” “burnt,” “moldy,” ...} for food). Search queries are then created by randomly juxtaposing an adjective a in either \mathbb{A}_D^+ or \mathbb{A}_D^- with a noun $n \in \mathbb{N}_D$, creating a set of appealing search queries \mathbb{Q}_D^+ and a set of unappealing search queries \mathbb{Q}_D^- respectively.

Using \mathbb{Q}_D^+ and \mathbb{Q}_D^- , we gather low-resolution thumbnails from stock-image websites, to form positive \mathbb{I}_D^+ and \mathbb{I}_D^- image sets respectively. Given the potential mismatch between the thumbnails and their search queries—particularly for unappealing images due to the scarcity of such content on stock websites—we select only the top matches from the search results to improve relevance.

To further refine the dataset and ensure its relevancy to the specific domain D , we implement a two-stage filtering process (Fig. 2). Initially, we use the BLIP [27] to produce a text description for each image, discarding any image whose description does not contain words related to D . This relevancy check is performed by comparing each word in the image’s description against WordNet [6] to find their lexical categories. Images are retained only if at least one word from the description is classified under a lexname that matches the domain. This approach minimizes the inclusion of irrelevant/out-of-domain images, enhancing the dataset’s quality and relevance to the application domain.

To ensure that objects from D occupy enough space in the image to be relevant (an image of a room with a small apple is not a food image), we assess this by generating a domain-relevancy map to identify and measure the extent of domain-related objects within an image as follows (Fig. 2).

The process starts by extracting noun phrases \mathbb{P} from the image’s text description [4]. From these, we identify phrases related to D as \mathbb{P}_D and for each $p \in \mathbb{P}_D$, we employ CLIPSeg [31] to segment the image I based on these phrases, resulting in a map $\text{CLIPSeg}(I, p)$ that highlights domain-relevant objects. By aggregating these maps and normalizing the combined map, we obtain the domain-relevancy map $M_D(I)$. Image are discarded when the pixel value sum of $M_D(I)$ is less than $\gamma \cdot w_I \cdot h_I$, where γ is a filtering threshold. To maintain dataset balance, we equalize the number of positive and negative images by removing excess ones from the larger subset.

Lastly, because all queried images are thumbnails and are fairly small (around 200×200), we apply ESRGAN [58] to upscale and zero-pad images to make them a reasonable size (512×512), making the final base image set \mathbb{I}_D .

3.2 Synthetic dataset image creation

One limitation of \mathbb{I}_D is the binary nature of our search queries yield images at the extreme ends of content appeal without capturing the subtle variations in between that are essential for training an accurate appeal score estimator.

To address this, we propose generating synthetic images with a nuanced spectrum of content appeal using a generative model like Stable Diffusion(SD) [47]. The process involves creating embeddings that encapsulate the characteristics

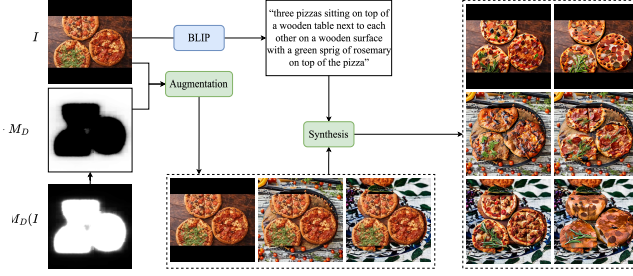


Fig. 3: Synthetic dataset creation. Given an image I , its text description, and its domain-relevancy map $M_D(I)$, we first locate “background” regions $1 - M_D(I)$ that should have minimal effect on content appeal. The image is first augmented using Stable Diffusion [47](Eq. (2)). We then use Textual Inversion [13] to generate appealing/unappealing-content embeddings, which can change image content appeal with respect to $M_D(I)$ (Eq. (1)).

of both appealing and unappealing content from our base image set, which can be linearly interpolated to produce all content appeal levels in between.

We start by selecting images in \mathbb{I}_D that best represent the highest and lowest appeal levels by gathering the top search results for search queries since early search results tend to be most relevant to queries themselves. Selected images $\mathbb{T}_D^+ \subset \mathbb{I}_D^+$ and $\mathbb{T}_D^- \subset \mathbb{I}_D^-$ are used to get “appealing” and “unappealing” embeddings z_D^+ and z_D^- with Textual Inversion [13], which capture the essence of content appeal at both ends of the spectrum. To represent any content appeal level between these two extremes we simply linearly blend these vectors $f(\alpha) = \alpha z_D^+ + (1 - \alpha) z_D^-$, where $\alpha \in [0, 1]$ controls the level of content appeal.

To adjust content appeal of an image I , we focus exclusively on the areas identified by the domain-relevancy map $M_D(I)$, leaving the background or non-domain content unchanged. This is crucial, especially in datasets where the subject’s appeal, such as food on a table or incidental items in room interiors, should not be influenced by their surroundings. While background contexts can affect perception, we initially set aside these influences for simplicity.

Specifically, we use the inpainting function of SD to adjust image content appeal. This function is denoted as $\text{SD}(I, p, M, \text{seed}())$, which takes the input image I and the text prompt p to change masked region M with randomization from the seed $\text{seed}()$. The adjusted image I' with the intended content appeal level $\alpha \in [0, 1]$ is produced through the equation:

$$I' = \text{SD}(I, \text{BLIP}(I) + f(\alpha), M_D(I), \text{seed}()). \quad (1)$$

To enrich our dataset with diverse background elements and mitigate the risk of overfitting our labeling algorithm, we employ the same inpainting function to freely generate any content for the background by using an empty string as the prompt, effectively allowing SD to introduce variability. The mask applied for this operation is the inverse of the domain-relevancy map $(1 - M_D(I))$, targeting




Fig. 4: Relative and absolute content appeal estimation. We use CLIP [43] image encoder, followed by several fully connected (FC) layers to predict the image relative content appeal difference (Fig. 4a) and absolute appeal (Fig. 4b).

the modification of non-domain areas in the image:

$$I'' = \text{SD}(I', \text{" "}, 1 - M_D(I), \text{seed}()), \quad (2)$$

where I' is the image that has been adjusted previously for content appeal.

While our procedural explanation first discusses adjusting content appeal and then background modification, the actual implementation first alters the background before the domain-specific content appeal. This sequence, depicted in our synthesis pipeline figure, is not expected to influence the outcome.

3.3 Relative content appeal estimation and final dataset annotation

To circumvent the laborious process of manually annotating a vast number of images, we propose employing an automatic labeling algorithm trained on synthetically generated data. This algorithm operates as a relative content appeal comparator, assessing the appeal difference between two images instead of determining an absolute appeal value, which demands a broader and more varied dataset for accurate estimation.

For this purpose, a synthetic dataset \mathbb{S} is created, producing N variations for each base image with differing content appeal levels α 's. and backgrounds, as previously outlined. We posit that the appeal difference between any two such variations correlates directly with the difference in their α values. Thus, for image I'_1 and I'_2 derived from the same original image I with content appeal parameters α_1 and α_2 , we assume their content appeal difference is $\hat{A}(I'_1, I'_2) = \alpha_1 - \alpha_2$.

A Siamese network architecture [25, 42], leveraging dual CLIP [43] image encoders with shared weights for feature extraction, serves as our relative content appeal comparator (Fig. 4a). This setup processes pairs of images, concatenates their features, and forwards these through fully connected layers to predict the appeal difference $\hat{A}_{pred}(I'_1, I'_2)$. The network is trained to minimize the discrepancy between predicted and assumed appeal differences $|\hat{A}_{pred}(I'_1, I'_2) - \hat{A}(I'_1, I'_2)|$.

After training, this comparator is tasked with labeling a comprehensive set of real images $\mathbb{I} = \mathbb{I}_D^+ \cup \mathbb{I}_D^-$ gathered from our initial queries, post domain-relevance filtering. To assign content appeal scores in the absence of absolute benchmarks,

we employ a voting mechanism using a subset of exemplar images $\mathbb{V}_D \subseteq \mathbb{I}$ as reference points. Each image’s appeal score is determined by averaging the comparator’s outcomes against these exemplars, subsequently scaling these scores to a 1-10 range to represent the spectrum of content appeal. This methodology facilitates the creation of our final ICAA dataset.

4 Absolute score estimation and enhancement

While the relative content appeal comparator has proven invaluable for dataset generation, our ultimate aim is to develop an estimator capable of assessing the content appeal of a single image in absolute terms. This leads to the creation of an absolute content appeal estimator, which evaluate the absolute content appeal of individual images $A(\cdot)$.

This estimator incorporates the same CLIP-based feature extraction mechanism used in the comparator, followed by a series of fully connected layers that culminate in a predicted content appeal score $A_{pred}(i)$. Training involves minimizing the discrepancy between the predicted appeal scores and the actual scores assigned during the dataset creation process $|A_{pred}(I) - A(I)|$.

An advantage of this absolute content appeal estimator is its utility in identifying and enhancing areas within an image that detract from its overall appeal (Fig. 1 Col. 3, 6). Instead of applying enhancement across the entire image, which risks altering already appealing or irrelevant regions, our goal is to specifically uplift areas deemed unappealing. A straightforward approach might involve applying a universal enhancement via Stable Diffusion, targeting maximum appeal. However, this method fails to discriminate between content that already meets or exceeds appeal thresholds and areas genuinely in need of improvement.

Although we can use $M_D(I)$ as a mask, this only resolves the first problem. So we generate a content appeal heatmap $M_D^H(I)$ that indicates the unappealing level at each image pixel (Fig. 5) to control the location and magnitude of enhancement. Given image I , we define a window w that slides over I to get overlapping image patches $w(I)$ every t pixels. The content appeal value of each pixel $p \in i$ is

$$\bar{A}(p) = \text{mean}_{p \in w(i)}(A(w(I))), \quad (3)$$

and the value of the image content appeal heatmap $M_D^H(I)$ for pixel p is $1 - n(\bar{A}(p))$, where $n(\cdot)$ normalizes all $\bar{A}(p)$ to the range $[0, 1]$. Lastly, we enhance the content appeal of I through

$$\text{SD}(I, \text{BLIP}(I) + z_D^+, M_D^H(I), \text{seed}()). \quad (4)$$

5 Experiments

5.1 Dataset creation

To show AID-AppEAL generalizes across different domains, we created two datasets with food and room interior images, both of which were automatically assigned appeal labels by our trained relative content appeal estimator:




Fig. 5: Image content appeal heatmap generation. We define a sliding window to capture overlapping patches of an image, where we use the content appeal estimator to estimate the content appeal score of each patch. The value of the heatmap for each pixel is averaged over all patches that include the pixel; we normalize all values and take their inverse, so a lighter color means the content in that region is more unappealing.

FOOD: Search queries were generated from the following sets of words:

- $\mathbb{N}_F = \{\text{"burger," "cake," "chicken," "cookie," "food," "rice," "pizza," "pasta," "salad," "steak," "yogurt"}\}$
- $\mathbb{A}_F^+ = \{\text{"delicious"}\}$
- $\mathbb{A}_F^- = \{\text{"burnt," "moldy," "rotten"}\},$

and use for [stock.adobe.com](https://www.stock.adobe.com) and [shutterstock.com](https://www.shutterstock.com). We generated 18,000 images for \mathbb{S}_F and 78,917 images for \mathbb{I}_F .

ROOM: Search queries were generated from the following sets of words:

- $\mathbb{N}_R = \{\text{"bathroom," "bedroom," "kitchen," "living room," "room"}\}$
- $\mathbb{A}_R^+ = \{\text{"interior"}\}$
- $\mathbb{A}_R^- = \{\text{"abandoned," "dirty"}\}.$

and generated 15,000 images for \mathbb{S}_R , 75,287 images for \mathbb{I}_R .

5.2 Model training

We train our relative content appeal comparator in two stages. In the first stage, we freeze the CLIP backbone and train the comparator on \mathbb{S}_D for 10 epochs using PyTorch’s AdamW optimizer with learning rate $1e^{-3}$ and batch size 16. In the second stage, we unfreeze the backbone and train the comparator for another 10 epochs with learning rate $1e^{-5}$. We then use the comparator to label \mathbb{I}_D and train the content appeal score estimator on it with the same training procedure as above. The final mean absolute error is $\text{MAE}_{i \in \mathbb{I}_D}(|A_{pred}(I) - A(I)|) = 0.6756$ for the FOOD dataset and 0.6332 for the ROOM dataset. We use Stable Diffusion v2.1 inpainting with depth-guided ControlNet [64] for appeal enhancement.

5.3 User study design

We conducted a user study to validate the effectiveness of our content appeal estimator and enhancer by comparing them against human preference. We invited 28 volunteers (male = 14, female = 14, non-binary = 0), with ages ranging 18 - 44, to participate in our study. After providing informed consent (IRB protocol number #anonymized), participants were asked to complete a survey on a computer screen in a lab setting.

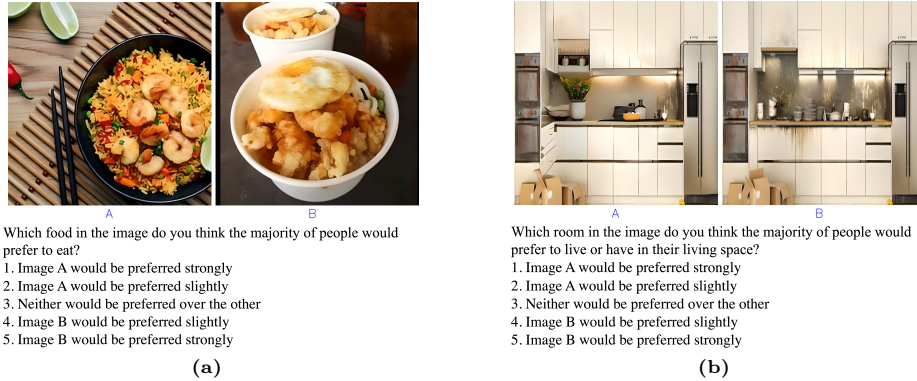


Fig. 6: User study interface. Participants are asked to answer questions by selecting one of the five options provided.

Table 1: Quantitative evaluation between content appeal labels and three IAA baselines. We evaluate the correlation coefficient between content appeal labels and three IAA baselines, and observe little to no correlation. RMSE metrics further suggests that our content appeal labels and IAA predictions are very different.

	FOOD			ROOM		
	DIAA	MPADA	NIMA	DIAA	MPADA	NIMA
PLCC	0.168	0.005	0.01	-0.123	-0.012	-0.147
SRCC	0.162	-0.015	0.003	-0.121	-0.017	-0.149
KRCC	0.109	-0.009	0.002	-0.082	-0.013	-0.098
RMSE	6.463	6.711	2.009	6.262	5.899	1.79

The survey presents each user with pairs of images, each of which accompanied by a question tailored to the specific image domain: “Which food in the image do you think the majority of the people would prefer **to eat**” (Fig. 6a) or “Which room in the image do you think the majority of the people would prefer **to live** or have in their living space” (Fig. 6b), which are phrased to focus attention on the subject matter rather than the image as a whole and direct responses towards a collective preference, thereby reducing the impact of personal tastes, which we verify through supplementary analysis to have minimal influence on the outcomes.

To mitigate potential biases linked to cultural or personal predispositions towards certain subjects, we ensured that each image pair featured the same kind of domain-relevant object (e.g., fried rice in Fig. 6a or a kitchen Fig. 6b). The study comprised two sections, one for food and the other for room images, with a mixture of real image pairs and pairs consisting of a real image alongside its enhanced version, randomly selected from our datasets to cover a broad spectrum of content appeal levels.

Participants are asked to answer a total of 74 questions: 38 in the food section and 36 in the room section, which included both comparisons between real




Fig. 7: Content appeal labels versus user preferences. Distribution of the difference in appeal labels $A(\text{image A}) - A(\text{image B})$ for each preference option in our user study (see Fig. 6). From response 1 to 5, we see a clear decrease in the mean of $A(\text{image A}) - A(\text{image B})$: as people start preferring B over A more, image B also becomes more appealing.




Fig. 8: Content appeal enhancement user responses. Percentage of responses for each category, where E represents the appeal-enhanced image, O is the original image, N is neither, and “pref” stands for “preferred.” We can see that 76.53% and 82.74% of the responses prefer the appeal enhanced images for the FOOD and ROOM dataset respectively.

images and assessments of enhancements. The presentation order of questions and the left/right positioning of images were randomized for each participant to prevent any ordering effects. Through this methodology, we collected 2072 individual responses, providing a comprehensive dataset for evaluating our systems against human judgment. Further details and the results of this evaluation are documented in our supplementary materials.

5.4 IAA baseline comparison

To show the difference between content appeal and image aesthetics, we uniformly stride every 1 out of every 100 images in each dataset and estimate their aesthetics scores with three popular, open-source IAA baselines [23, 53, 57]. We then compare to our appeal labels and observe little correlation as well as large value difference between the two (Tab. 1). Refer to the supplementary for more details.



Fig. 9: Food image content appeal enhancement comparison with baselines. Compared with baselines, our enhancer respects the color, texture, and structures of the original images while effectively improving the appeal level of their content.

5.5 Human preference comparison

Out of 2072 responses, 672 of them compared the appeal of real images in the FOOD and ROOM datasets, respectively. We first assess the accuracy of our labeling process by plotting in Fig. 7 the distribution of appeal label differences, $A(\text{image A}) - A(\text{image B})$, versus user preference, and observe a clear decrease in mean values from response 1 to 5. This indicates that image B’s appeal increases relative to image A aligns with user preferences and demonstrates that our content appeal labels are indeed accurate.

To compare content appeal before and after appeal enhancement (Fig. 8), we received 392 responses for the FOOD dataset and 336 for the Room dataset out of the 2072 responses. For the former, 76.53% of responses favored enhanced images, with 41.58% strongly preferring them. In the ROOM dataset, 82.74% preferred enhanced images, with 53.87% showing strong preference. This demonstrates a clear preference for enhanced images from our methods across both datasets.

5.6 Content appeal enhancer baseline comparison

We have chosen existing diffusion-based image editing methods InstructPix2Pix (IP2P) [5], Null-text Inversion (NTI) [35], pix2pix-zero (P2P0) [39], and Text2LIVE (T2L) [3], which we think can be applied in a similar setting to ours.

In Figs. 9 and 10, we provide a visual comparison between our method and these baselines, as well as our content appeal heatmaps M_D^H . As we can see,

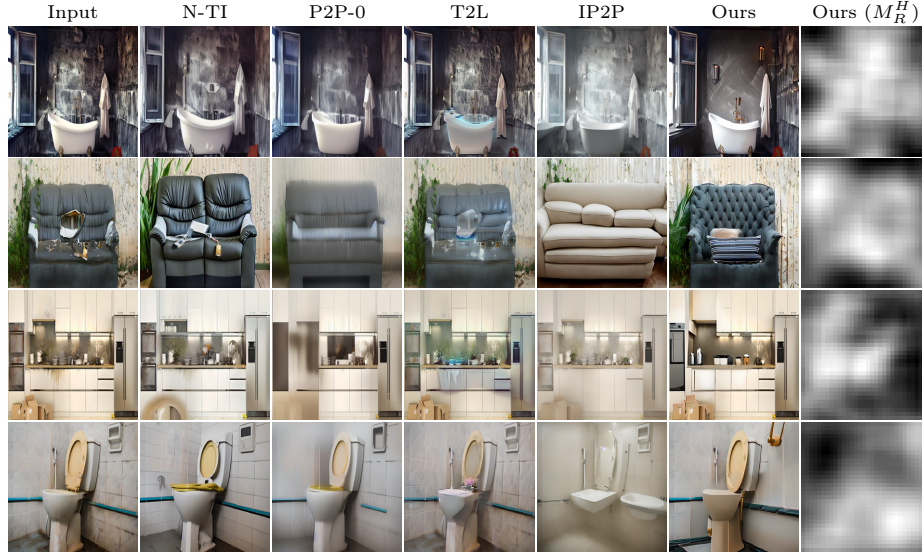


Fig. 10: Room image content appeal enhancement comparison with baselines. Our enhancer better respects the color, texture, and structures of the original images while effectively improving the appeal level of their content than baselines.

N-TI tends to enlarge objects in images (Fig. 9, Rows 1-3) without changing the content appeal level much. It may also produces drastic undesired change to the change (Fig. 9, Row 4; Fig. 10, Row 2). P2P-0 and T2L often blur objects and create shadowing artifacts with little effect on the content appeal; IP2P is prone to changing the images too much (Fig. 9, Rows 3-4; Fig. 10, Rows 1, 2, 4). In contrast, our method is able to constrain the location and the magnitude of image content appeal enhancement using the heatmap and produce results with improved content appeal while respecting the color and structure of the input images. Please refer to the supplementary for baseline details, more results, and the ablation study.

6 Conclusion

In this work, we explored a new area of image appeal assessment (ICAA) that evaluates the interest an image creates in observers. We highlight the challenge of manual labeling in dataset creation, and propose a fully automated pipeline to generate extensive datasets across domains. Our research illustrates how these datasets can be used to train an appeal estimator and facilitate appeal enhancement applications. Validation of our methods is conducted through a user study, confirming their effectiveness.

References

1. Avrahami, O., Lischinski, D., Fried, O.: Blended diffusion for text-driven editing of natural images. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 18208–18218 (2022)
2. Bansal, H., Grover, A.: Leaving reality to imagination: Robust classification via generated datasets. ArXiv **abs/2302.02503** (2023)
3. Bar-Tal, O., Ofri-Amar, D., Fridman, R., Kasten, Y., Dekel, T.: Text2LIVE: Text-driven layered image and video editing. In: European Conference on Computer Vision (ECCV). pp. 707–723. Springer (2022)
4. Bird, S., Klein, E., Loper, E.: Natural language processing with Python: analyzing text with the natural language toolkit. " O'Reilly Media, Inc." (2009)
5. Brooks, T., Holynski, A., Efros, A.A.: InstructPix2Pix: Learning to follow image editing instructions. arXiv preprint arXiv:2211.09800 (2022)
6. Brown, K.: Encyclopedia of language and linguistics, vol. 1. Elsevier (2005)
7. Carlini, N., Hayes, J., Nasr, M., Jagielski, M., Sehwal, V., Tramèr, F., Balle, B., Ippolito, D., Wallace, E.: Extracting training data from diffusion models. ArXiv **abs/2301.13188** (2023)
8. Chang, K.Y., Lu, K.H., Chen, C.S.: Aesthetic critiques generation for photos. In: International Conference on Computer Vision (ICCV). pp. 3514–3523 (2017)
9. Datta, R., Li, J., Wang, J.Z.: Algorithmic inferencing of aesthetics and emotion in natural images: An exposition. In: IEEE International Conference on Image Processing (ICIP). pp. 105–108. IEEE (2008)
10. Dhar, S., Ordonez, V., Berg, T.L.: High level describable attributes for predicting aesthetics and interestingness. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1657–1664. IEEE (2011)
11. Dhariwal, P., Nichol, A.: Diffusion models beat GANs on image synthesis. Advances in Neural Information Processing Systems (NeurIPS) **34**, 8780–8794 (2021)
12. Fang, Y., Zhu, H., Zeng, Y., Ma, K., Wang, Z.: Perceptual quality assessment of smartphone photography. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3677–3686 (2020)
13. Gal, R., Alaluf, Y., Atzmon, Y., Patashnik, O., Bermano, A.H., Chechik, G., Cohen-Or, D.: An image is worth one word: Personalizing text-to-image generation using textual inversion (2022). <https://doi.org/10.48550/ARXIV.2208.01618>, <https://arxiv.org/abs/2208.01618>
14. Ghadiyaram, D., Bovik, A.C.: Massive online crowdsourced study of subjective and objective picture quality. IEEE Transactions on Image Processing (TIP) **25**(1), 372–387 (2015)
15. Göring, S., Raake, A.: Image appeal revisited: Analysis, new dataset and prediction models. IEEE Access (2023)
16. He, S., Zhang, Y., Xie, R., Jiang, D., Ming, A.: Rethinking image aesthetics assessment: Models, datasets and benchmarks. In: International Joint Conference on Artificial Intelligence (IJCAI) (2022)
17. He, X., Wandt, B., Rhodin, H.: LatentKeypointGAN: Controlling gans via latent keypoints. arXiv preprint arXiv:2103.15812 (2021)
18. Hosu, V., Goldlucke, B., Saupe, D.: Effective aesthetics prediction with multi-level spatially pooled features. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 9375–9383 (2019)
19. Hosu, V., Lin, H., Sziranyi, T., Saupe, D.: KonIQ-10K: An ecologically valid database for deep learning of blind image quality assessment. IEEE Transactions on Image Processing (TIP) **29**, 4041–4056 (2020)

20. Jinjin, G., Haoming, C., Haoyu, C., Xiaoxing, Y., Ren, J.S., Chao, D.: PIPAL: A large-scale image quality assessment dataset for perceptual image restoration. In: European Conference on Computer Vision (ECCV). pp. 633–651. Springer (2020)
21. Ke, J., Wang, Q., Wang, Y., Milanfar, P., Yang, F.: MUSIQ: Multi-scale image quality transformer. In: International Conference on Computer Vision (ICCV). pp. 5148–5157 (2021)
22. Kong, S., Shen, X., Lin, Z., Mech, R., Fowlkes, C.: Photo aesthetics ranking network with attributes and content adaptation. In: European Conference on Computer Vision (ECCV). pp. 662–679. Springer (2016)
23. Kong, S., Shen, X., Lin, Z., Mech, R., Fowlkes, C.: Photo aesthetics ranking network with attributes and content adaptation. In: European Conference on Computer Vision (ECCV) (2016)
24. Lao, S., Gong, Y., Shi, S., Yang, S., Wu, T., Wang, J., Xia, W., Yang, Y.: Attentions help cnns see better: Attention-based hybrid image quality assessment network. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1140–1149 (2022)
25. Lee, J.T., Kim, C.S.: Image aesthetic assessment based on pairwise comparison - a unified approach to score regression, binary classification, and personalization. International Conference on Computer Vision (ICCV) pp. 1191–1200 (2019)
26. Li, D., Jiang, T., Jiang, M.: Norm-in-norm loss with faster convergence and better performance for image quality assessment. In: ACM International Conference on Multimedia (ACMMM). pp. 789–797 (2020)
27. Li, J., Li, D., Xiong, C., Hoi, S.: BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In: International Conference on Machine Learning (ICML). pp. 12888–12900. PMLR (2022)
28. Lin, H., Hosu, V., Saupe, D.: KADID-10K: A large-scale artificially distorted IQA database. In: International Conference on Quality of Multimedia Experience (QoMEX). pp. 1–3. IEEE (2019)
29. Liu, B., Zhu, Y., Song, K., Elgammal, A.: Towards faster and stabilized GAN training for high-fidelity few-shot image synthesis. In: International Conference on Learning Representations (ICLR) (2021)
30. Lu, X., Lin, Z., Jin, H., Yang, J., Wang, J.Z.: RAPID: Rating pictorial aesthetics using deep learning. In: ACM International Conference on Multimedia (ACMMM). pp. 457–466 (2014)
31. Lüddecke, T., Ecker, A.: Image segmentation using text and image prompts. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7086–7096 (June 2022)
32. Ma, K., Duanmu, Z., Wu, Q., Wang, Z., Yong, H., Li, H., Zhang, L.: Waterloo exploration database: New challenges for image quality assessment models. IEEE Transactions on Image Processing (TIP) **26**(2), 1004–1016 (2016)
33. Ma, S., Liu, J., Wen Chen, C.: A-Lamp: Adaptive layout-aware multi-patch deep convolutional neural network for photo aesthetic assessment. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4535–4544 (2017)
34. Meng, C., He, Y., Song, Y., Song, J., Wu, J., Zhu, J.Y., Ermon, S.: SDEdit: Guided image synthesis and editing with stochastic differential equations. In: International Conference on Learning Representations (ICLR) (2022)
35. Mokady, R., Hertz, A., Aberman, K., Pritch, Y., Cohen-Or, D.: Null-text inversion for editing real images using guided diffusion models. arXiv preprint arXiv:2211.09794 (2022)

36. Murray, N., Marchesotti, L., Perronnin, F.: AVA: A large-scale database for aesthetic visual analysis. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 2408–2415. IEEE (2012)
37. Nafchi, H.Z., Shahkolaei, A., Hedjam, R., Cheriet, M.: Mean deviation similarity index: Efficient and reliable full-reference image quality evaluator. *Ieee Access* **4**, 5579–5590 (2016)
38. Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: GLIDE: Towards photorealistic image generation and editing with text-guided diffusion models. arXiv preprint arXiv:2112.10741 (2021)
39. Parmar, G., Singh, K.K., Zhang, R., Li, Y., Lu, J., Zhu, J.Y.: Zero-shot image-to-image translation. arXiv preprint arXiv:2302.03027 (2023)
40. Ponomarenko, N., Ieremeiev, O., Lukin, V., Egiazarian, K., Jin, L., Astola, J., Vozel, B., Chehdi, K., Carli, M., Battisti, F., et al.: Color image database TID2013: Peculiarities and preliminary results. In: European Workshop on Visual Information Processing (EUVIP). pp. 106–111. IEEE (2013)
41. Ponomarenko, N., Lukin, V., Zelensky, A., Egiazarian, K., Carli, M., Battisti, F.: A database for evaluation of full-reference visual quality assessment metrics. *Advances of Modern Radioelectronics* **10**(4), 30–45 (2009)
42. Prashnani, E., Cai, H., Mostofi, Y., Sen, P.: PieAPP: Perceptual image-error assessment through pairwise preference. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1808–1817 (2018)
43. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning (ICML). pp. 8748–8763. PMLR (2021)
44. Ramesh, A., Dhariwal, P., Nichol, A., Chu, C., Chen, M.: Hierarchical text-conditional image generation with CLIP latents. ArXiv **abs/2204.06125** (2022)
45. Redi, J., Hossfeld, T., Korshunov, P., Mazza, F., Povoas, I., Keimel, C.: Crowdsourcing-based multimedia subjective evaluations: a case study on image recognizability and aesthetic appeal. In: ACM Workshop on Crowdsourcing for Multimedia (CrowdMM) (2013)
46. Ren, J., Shen, X., Lin, Z., Mech, R., Foran, D.J.: Personalized image aesthetics. In: International Conference on Computer Vision (ICCV) (Oct 2017)
47. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models (2021)
48. Sauer, A., Chitta, K., Müller, J., Geiger, A.: Projected GANs converge faster. *Advances in Neural Information Processing Systems (NeurIPS)* **34**, 17480–17492 (2021)
49. Sauer, A., Karras, T., Laine, S., Geiger, A., Aila, T.: StyleGAN-T: Unlocking the power of GANs for fast large-scale text-to-image synthesis. arXiv preprint arXiv:2301.09515 (2023)
50. Sauer, A., Schwarz, K., Geiger, A.: StyleGAN-XL: Scaling stylegan to large diverse datasets. In: ACM Special Interest Group on Computer Graphics and Interactive Techniques (SIGGRAPH). pp. 1–10 (2022)
51. Savakis, A.E., Etz, S.P., Loui, A.C.: Evaluation of image appeal in consumer photography. In: Human Vision and Electronic Imaging (HVEI). vol. 3959, pp. 111–120. SPIE (2000)
52. Sheikh, H.R., Sabir, M.F., Bovik, A.C.: A statistical evaluation of recent full reference image quality assessment algorithms. *IEEE Transactions on Image Processing (TIP)* **15**(11), 3440–3451 (2006)

53. Sheng, K., Dong, W., Ma, C., Mei, X., Huang, F., Hu, B.G.: Attention-based multi-patch aggregation for image aesthetic assessment. In: ACM International Conference on Multimedia (ACMMM). pp. 879–886 (2018)
54. Shrivastava, A., Pfister, T., Tuzel, O., Susskind, J., Wang, W., Webb, R.: Learning from simulated and unsupervised images through adversarial training. In: Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017)
55. Siahaan, E., Hanjalic, A., Redi, J.: A reliable methodology to collect ground truth data of image aesthetic appeal. *IEEE Transactions on Multimedia (TMM)* **18**, 1338–1350 (2016)
56. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (ICLR). OpenReview.net (2021), <https://openreview.net/forum?id=St1giarCHLP>
57. Talebi, H., Milanfar, P.: NIMA: Neural image assessment. *IEEE Transactions on Image Processing (TIP)* **27**(8), 3998–4011 (2018)
58. Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., Change Loy, C.: ES-RGAN: Enhanced super-resolution generative adversarial networks. In: European Conference on Computer Vision (ECCV). pp. 0–0 (2018)
59. Wang, Z., Zheng, H., He, P., Chen, W., Zhou, M.: Diffusion-GAN: Training GANs with diffusion. *arXiv preprint arXiv:2206.02262* (2022)
60. Xue, W., Zhang, L., Mou, X., Bovik, A.C.: Gradient magnitude similarity deviation: A highly efficient perceptual image quality index. *IEEE Transactions on Image Processing (TIP)* **23**(2), 684–695 (2013)
61. Yang, Y., Xu, L., Li, L., Qie, N., Li, Y., Zhang, P., Guo, Y.: Personalized image aesthetics assessment with rich attributes. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 19861–19869 (2022)
62. Ying, Z., Niu, H., Gupta, P., Mahajan, D., Ghadiyaram, D., Bovik, A.: From patches to pictures (PaQ-2-PiQ): Mapping the perceptual space of picture quality. In: Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3575–3585 (2020)
63. Zhang, B., Niu, L., Zhang, L.: Image composition assessment with saliency-augmented multi-pattern pooling. *arXiv preprint arXiv:2104.03133* (2021)
64. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: International Conference on Computer Vision (ICCV). pp. 3836–3847 (2023)
65. Zhu, H., Li, L., Wu, J., Zhao, S., Ding, G., Shi, G.: Personalized image aesthetics assessment via meta-learning with bilevel gradient optimization. *IEEE Transactions on Cybernetics (TCYB)* (2020), doi:10.1109/TCYB.2020.2984670