

# Detecting Cultural Differences in News Video Thumbnails via Computational Aesthetics

Marvin Limpijankit<sup>1</sup>, John R. Kender<sup>1</sup>

<sup>1</sup>Columbia University  
New York, NY 10027

ml4431@columbia.edu, jrk@cs.columbia.edu

## Abstract

We propose a two-step approach for detecting differences in the style of images across sources of differing cultural affinity, where images are first clustered into finer visual themes based on content before their aesthetic features are compared. We test this approach on 2,400 YouTube video thumbnails taken equally from two U.S. and two Chinese YouTube channels, and relating equally to COVID-19 and the Ukraine conflict. Our results suggest that while Chinese thumbnails are less formal and more candid, U.S. channels tend to use more deliberate, proper photographs as thumbnails. In particular, U.S. thumbnails are less colorful, more saturated, darker, more finely detailed, less symmetric, sparser, less varied, and more up close and personal than Chinese thumbnails. We suggest that most of these differences reflect cultural preferences, and that our methods and observations can serve as a baseline against which suspected visual propaganda can be computed and compared.

## 1 Introduction

Visual imagery is an extremely important channel for international news coverage, as it connects the audience to events that may otherwise be removed from direct experience, and serves as a powerful tool for conveying emotion (Smith and Joffe 2013). For threatening events, such as global warming, wars, or pandemics, images have also been shown to be significant drivers of public engagement, “providing viewers with tangible and emotion-evoking examples that act as visual proof” (Joffe 2008; Martikainen and Sakki 2021). As such, highly visual mediums like YouTube play a significant role in modern news consumption (Burgess and Green 2018; Xie et al. 2011).

In this paper, we apply computational approaches to quantify differences in the use of visual imagery between news sources from varying regional backgrounds. More specifically, our work lies at the interplay of sociology and computer science, adopting techniques from computer vision to extract visual features from YouTube thumbnails at scale, before using methods such as dimensionality reduction to accentuate cross-cultural differences between U.S. and Chinese channels. We focus on two international events, the COVID-19 pandemic, and the war in Ukraine.

Since statistical image properties are highly dependent on the content of the image (Kao, He, and Huang 2017; Machajdik and Hanbury 2010), our proposed framework first clusters images about a given news story by “visual themes” (i.e. different aspects or sub-events), We then analyze the aesthetic differences between channels. Lastly, we consider the views, likes, and comments for each video, in order to capture potential cross-cultural differences in how thumbnail properties may influence video performance.

Our contributions in this paper are:

1. A framework to compare image properties in YouTube thumbnails that accounts for visual thematic content.
2. An algorithmic extraction of 21 visual aesthetics of a dataset of 2,400 thumbnail images, taken from two U.S. and two Chinese channels, covering two separate international events, COVID-19 and the Ukraine war.
3. An analysis of the culture-specific differences in the use of these aesthetics, and their relation to viewership statistics and potential disinformation.

## 2 Related Work

### 2.1 Computational Aesthetics

In automatic image aesthetic assessment, also known as computational aesthetics, a model is trained to quantify the beauty of an image through some aesthetic score. Often, these models assess an image by measuring the degree to which it adheres to fundamental but abstract principles of photography, such as balance, rhythm, harmony, unity, tone, etc. (Anwar et al. 2021).

There are two main approaches to image aesthetic assessment: hand-crafted methods, and deep learning methods.

**Hand-crafted methods.** These rely on extracting pre-defined low-level features from images and fitting a machine learning model, such as SVM, to do the prediction. Examples of features that have been extensively explored previously include sharpness, colorfulness, contrast, and texture (Lo, Liu, and Chen 2013; Redi et al. 2015; Aydın, Smolic, and Gross 2014).

**Deep learning methods.** These leverage neural networks to learn low and high-level image features without (or with relatively little) human instruction. These are trained on large image datasets such as Photo.net (Datta et al.

2006), AVA (Murray, Marchesotti, and Perronnin 2012), and CUHK-PQ (Luo, Wang, and Tang 2011), and the specific architecture can vary. Doshi et al. (Doshi, Shikkenawis, and Mitra 2020) fine-tuned two pre-trained models (AlexNet and VGG) to classify images into high and low aesthetic categories. Kao et al. (Kao, He, and Huang 2017) investigated using a Multi-Task Convolutional Neural network (MTCNN) to jointly perform aesthetic assessment and semantic recognition.

**Applications.** In addition to predicting image aesthetics such as quality, beauty, and liking (Bartho, Thoemmes, and Redies 2023; Peng and JEMMOTT III 2018), there is evidence to suggest that computational aesthetics can serve as a valuable starting point for a variety of downstream tasks.

For instance, one experiment demonstrated the ability to infer the personality traits (“Big Five Inventory”) of Flickr users based on the color, composition, texture, and face features of images that the user tagged as ‘favorites’ on the platform (Segalin et al. 2016).

Likewise, Machajdik et al. (Machajdik and Hanbury 2010) achieved state-of-the-art results on image emotion classification using a similar set of features, including color, texture, composition, and content, inspired by psychology and art theory. However, results were not uniform: content features such as the number of faces and relative size of the biggest face were good discriminators between fear/disgust vs. amusement in the International Affective Picture System dataset, whereas color features were more predictive in an artistic photography dataset collected from the DeviantArt website.

These studies reflect the ability of image aesthetics to emphasize certain feelings, but warn of the importance of considering the content and source.

## 2.2 Image Analysis in News Media

Cross-cultural differences in news portrayals have been studied extensively from both a sociological (manual) and computational (automatic) perspective.

**Manual.** One investigation analyzed images of the COVID-19 pandemic appearing in Finnish newspapers, focusing on the social representations of individuals from varying age groups (Martikainen and Sakki 2021). The authors noted that the angle of the photo often contributed to visual rhetoric, for instance, images of children studying at home emphasized a high angle, placing the spectator in the authoritative position overseeing the child whereas those of young adults tended to be taken at eye-level.

Similarly, in a study on how the media covered Pope Francis’ visit to Cuba, Thomson et al. discovered that photo angles in U.S. sources framed the Pope at higher positions than Cuban politicians whereas local sources used a more equal leveling (Thomson, Perreault, and Duffy 2018). Another contrast was the content of the images, with Cuban media focusing almost exclusively on the Pope, whereas U.S. media more often published photographs of Cuban people. (Fahmy 2010; Rafiee, Spooen, and Sanders 2023).

**Computational.** Although these frequently involve analyzing text data to compare sentiments and topics (Wallbing 2021; Chen et al. 2023), the visual medium of news media has also been explored. Bhargava et al. (Bhargava, Bishop, and Zuckerman 2020) proposed Photo Spaces, a visualization tool to describe visual narratives for an event, based on image embeddings produced by a pre-trained ResNet-50 model. The spatial positioning of images in the photo-space is determined by reducing these embeddings into two dimensions, via the Uniform Manifold Approximation and Projection technique, from which clusters corresponding to visual themes can be extracted. The authors document an example of the tool (See Figure 1) on abortion rights as a topic, adding the political affinity of sources (right- vs. left-leaning) through the color of each image’s border.



Figure 1: An example of Photo-Space visualization (Bhargava, Bishop, and Zuckerman 2020), where physical placement captures thematic relations, and image borders are color-coded by political affinity.

Similarly, Zhang et al. (Zhang and Peng 2022) evaluated the effectiveness of bag-of-visual-words, self-supervised learning, and transfer learning to cluster unlabeled image data. They conclude that transfer learning using pre-trained models to obtain embeddings outperforms the other methods, though the choice of the pre-trained model significantly affected clustering results.

Following these trends, in our investigation we adopt image clustering as a pre-processing step in our pipeline, as a way of controlling for visual context.

## 2.3 Video Thumbnails

Thumbnails are small image representations of full-length videos. Thumbnails “tell the users what the video is about (i.e., be informative), and grab potential viewers’ attention (i.e., be visually appealing)” (Koh and Cui 2022).

Kim et al. (Kim et al. 2023) evaluated the design of thumbnails for data-driven news articles, concluding similarly that for thumbnails involving informative graphics, users prefer ones that are interpretable and attractive. For branded content, a thumbnail’s success (measured in view-through) can also be attributed to visual features, such as colorfulness and brightness (Koh and Cui 2022).

As such, our work pursues the observations that thumbnails may serve as an interesting visual medium for cross-cultural news analysis, reflecting both the content the media

wishes to present (information) as well as their attempts to appeal to specific cultural perspectives (attraction).

### 3 Methods

#### 3.1 Data Collection

**Choice of channel.** To evaluate cross-cultural differences, we selected two U.S. and two Chinese YouTube channels. The U.S. sources are ABC and CBS, with 15.6M and 5.5M subscribers respectively, and the Chinese sources are CGTN (Chinese Global Television Network) and New China TV, with 3.0M and 1.4M subscribers respectively.

**Choice of event.** Inspired by previous work (Martikainen and Sakki 2021), we selected the COVID-19 pandemic as one of the international news events to examine. The complexity of the event gives rise to a diverse set of event themes—medical, economic, political—which may transfer to detectable visual themes. For similar reasons, we also selected the war in Ukraine, a more recent event with different event themes—military conflict, the humanitarian crisis, and geopolitics.

**Thumbnail curation.** The YouTube API was used to scrape the video metadata using a search query (“covid 19”, “ukraine war”), the channel\_ids of the above, and a published\_after date. For each event and each channel, we collected the top 300 videos, sorted by YouTube’s default “relevance” parameter. Iterating through the results, image content is then scraped using a simple Python script, and pre-processed to remove any vertical and horizontal black bars using the Pillow package. Each of the resulting 2,400 thumbnails were recorded, associated with an image\_id, channel, view count, like count, and comment data.

#### 3.2 Visual Theme Identification

Prior to any comparative analysis, in order to control for image content, the thumbnails are first divided into subgroups based on visual themes.

**Thumbnail clustering.** After some experimentation, our best results were obtained by using Concept (Grootendorst 2021), a CLIP and BERTopic-based open source tool that performs topic modeling on images, by first encoding images using CLIP, and then identifying clusters on top of those embeddings. An example of the results produced using this method is shown in Figure 2.

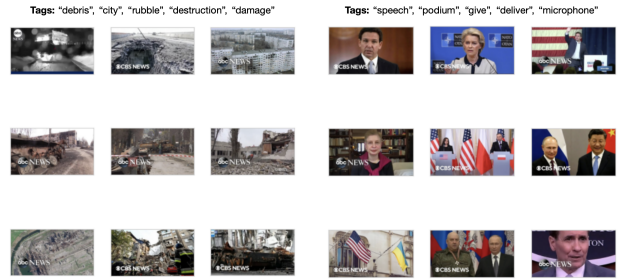


Figure 2: Example visual themes from the Ukraine war.

**Theme tagging.** Thumbnail clusters themselves carry no easily interpretable meaning, so to understand the visual themes being extracted we associated each cluster with five textual tags. We leveraged Tag2Text (Huang et al. 2023), a model capable of recognizing 3,429 commonly human-used categories, to generate tags for every image. Then, on a cluster level, we aggregated these tags to generate five tags that are most representative of the cluster’s theme.

For this, we examined three approaches:

1. Use the top five most commonly occurring tags.
2. Use the same approach as the first but with the added requirement that these tags appear less than some  $k$  times across all images, including other clusters.
3. Compute a tf-idf (term frequency-inverse document frequency) score for each tag with respect to a cluster using Equation 1, and use the five tags with the largest weights.

$$w_{i,j} = tf_{i,j} \log \left( \frac{N}{df_i} \right) \quad (1)$$

where  $w_{i,j}$  is the tf-idf weight of tag  $i$  in cluster  $j$ ,  $tf_{i,j}$  is the frequency of the tag in the cluster,  $N$  is the number of clusters, and  $df_i$  is the number of clusters containing the tag.

The second and third approaches were introduced to filter out certain tags, such as man, that appear in almost every thumbnail and thus are not discriminative between clusters. The results from these tagging methods respectively are summarized in Table 1. We selected the second approach, which are charted in Figure 3.

Those four Ukraine themes, roughly “speech”, “interview”, “rubble”, and “army”, were distributed in a 2:4:1:3 ratio overall, heavily favoring talk over destruction, and with Chinese thumbnails even more so.

Those three COVID themes, roughly “speech”, “interview”, and “equipment”, were distributed in a 1:1:4 ratio overall, heavily favoring “equipment” by both cultures. But Chinese thumbnails favored “speech” over “interview”; U.S. thumbnails were the reverse.

#### 3.3 Image Feature Extraction

Image features are extracted using a hand-crafted approach, where various pre-defined statistical measures are computed for each image. An exhaustive reference list of the 21 features, in eight feature groups, is displayed in Table 2. The majority of our feature set is adopted from Bartho et

Event	Theme	Tagging Method		
		Approach 1 (top 5)	Approach 2 (top 5 adjusted)	Approach 3 (tf-idf)
COVID	0	woman, news, person, stand, man	interview, sit, table, blue, speech	mail, tie, interview, sunglasses, give
	1	news, wear, tie, business suit, man	business suit, tie, suit, interview, microphone	bookshelf, office, interview, black, tie
	2	stand, mask, man, wear, person	mask, equipment, syringe, hand, liquid	worker, tablet, syringe, hospital room, garment
Ukraine war	0	business suit, speech, stand, tie, man	speech, podium, give, deliver, microphone	give, podium, business suit, speech, tie
	1	business suit, wear, tie, stand, man	business suit, news, interview, woman, suit	shake, suit, interview, business suit, tie
	2	man, rubble, debris, city, building	debris, city, rubble, destruction, damage	explosion, flame, damage, garbage, rubble
	3	stand, army, camouflage, person, man	camouflage, army, soldier, gun, equipment	weapon, load, gun, soldier, rifle

Table 1: Methods for theme tagging; we chose the second.




Figure 3: Ukraine and COVID themes, China vs. U.S.

al. (Bartho, Thoemmes, and Redies 2023). A brief justification of each feature is included here, but the reader is encouraged to consult the references for the computational details and tuning parameters of each algorithm.

**Color.** *Hue, Saturation, Lab(a), Lab(b), Color Entropy.* The use of colors can affect the perceived mood of an image. Although there are multiple ways to represent colors numerically, the HSV space (for hue, saturation, and value), and the  $L^*a^*b^*$  space (for lightness, green-red, and blue-yellow) both have been shown to be more aligned with the human perception of color (Fadaei 2021). Color-entropy captures the visual percept of the “colorfulness” of an image.

**Dimension.** *Aspect Ratio, Image Size.* Although thumbnails are reshaped into pre-defined dimensions by YouTube, videos vs. shorts are presented slightly differently. Channels may also choose to include border padding within the thumbnail image itself, adding variation to dimension features since this is cropped in pre-processing.

**Lightness.** *Contrast, Luminance, Luminance Entropy.* To complement color, several features for lightness (i.e. brightness) are also considered. For instance, the effect of a very bright vibrant red such as Carmine is not be equivalent to that of a darker red such as Burgundy.

**HOG.** *Self-Similarity, Complexity, Anisotropy.* A Histogram of Oriented Gradients (HOG) is a feature descriptor that extracts texture information, originally introduced for object recognition. The method captures statistics about the distribution of the angles of color edges. It first establishes a histogram whose bins span a quantized set of orientations, and then accumulates within each bin the  $L^*a^*b^*$  gradient magnitudes that occur at that angle. Each histogram is local to a single cell of a mesh-like partition of an image, and it therefore captures the edge activity within a local region. These orientation statistics are then used to determine approximations to three more subjectively meaningful visual percepts.

*Self-similarity* measures the overall likeness of the HOG

features of a localized region of an image compared to its neighboring regions. It is a useful feature for separating visually noisy objects from plain or even complicated backgrounds.

*Complexity* is calculated as the average color gradient strength throughout an image. An image with many sharp color edges, and therefore very strong gradients, has high complexity; cartoons have low complexity.

*Anisotropy* measures the relative strengths of gradients at different orientations, calculated as the standard deviation of HOG histogram bin magnitudes. High values would suggest strongly differing strengths at different orientations, and are typical of human-created objects with parallel and perpendicular edges.

**Fourier.** *Fourier Slope, Fourier Sigma.* The Fourier Transform maps an image to the frequency domain, decomposing the image into a sum of two-dimensional sine and cosine components. Low-frequency components capture less variant parts of the image, whereas high-frequency components capture sharp details. The structure of the image as a whole can be represented by charting how rapidly the logarithm of the power (magnitude squared) of the Fourier transformation decreases from low-frequency components to high-frequency components.

*Fourier slope* is calculated as the slope of the line of best fit, whereas the *Fourier sigma* is calculated as the root mean squared error of that line. An image with larger (less negative) slope suggests a stronger presence of details, often natural, whereas smaller values suggest an overall smoother image, often man-made.

**Symmetry.** *Symmetry-LR, Symmetry-UD.* Pre-trained Convolutional Neural Network (CNN) models, such as AlexNet, generally learn low-level features such as color and texture in their initial convolutional layers. These activation layers can be used to produce a symmetry metric that is closely aligned with human perception. The computation simply compares their values against their flipped versions: a horizontal flip for *lr* (left-right) symmetry (see Figure 4), and a vertical flip for *ud* (up-down) symmetry.

**CNN.** *Sparseness, Variability.* The activation layers of a CNN, such as AlexNet, can also be used to calculate higher level human percepts.

*Sparseness* can be defined by computing the variances in a CNN’s initial response maps. Images with high sparseness often correspond to images with homogeneous patches, such as images with only a few objects; see Figure 4.

*Variability* can be measured as the variance over all response maps, and captures something like the inverse of *self-similarity*, such as with vegetation imagery.

**Other.** *Shot Scale, Bag-of-objects.* we complemented aesthetic features with two additional features reflecting photographic composition.

*Shot scale.* From the CineScale project, we used a trained CNN to predict the shot scale of an image as close, medium, or long; see Figure 4.

*Bag-of-objects.* Using the CNN YOLOv5, we represented each image as an object histogram, where each bin

Group	Literature Reference(s)	Dimensions	Feature	Description
Color	(Datta et al. 2006), (Li and Chen 2009), (Mallon, Redies, and Hayn-Leichsenring 2014), (Schifanella, Redi, and Aiello 2015), (Thieleking et al. 2020), (Iigaya et al. 2021), (Geller et al. 2022)	5	<i>Hue</i>	mean of hue channel in HSV space
			<i>Saturation</i>	mean of saturation channel in HSV space
			<i>Lab(a)</i>	mean of a* channel in L*a*b* space
			<i>Lab(b)</i>	mean of b* channel in L*a*b* space
Dimension	(Datta et al. 2006), (Iigaya et al. 2021) (Mallon, Redies, and Hayn-Leichsenring 2014)	2	<i>Color Entropy</i>	Shannon entropy of hue channel in HSV space
			<i>Aspect Ratio</i>	image width divided by image height (pixels)
Lightness	(Schifanella, Redi, and Aiello 2015), (Sidhu et al. 2018), (Mather 2018), (Iigaya et al. 2021), (Peli 1990)	3	<i>Image Size</i>	image width + image height (pixels)
			<i>Contrast</i>	standard deviation of L* channel in L*a*b* space
HOG	(Bosch, Zisserman, and Munoz 2007), (Braun et al. 2013)	3	<i>Luminance</i>	mean of L* channel in L*a*b* space
			<i>Luminance Entropy</i>	Shannon entropy of L* channel in L*a*b* space
			<i>Self-Similarity</i>	similarity of HOG features
Fourier	(Redies et al. 2008)	2	<i>Complexity</i>	mean gradient strength
			<i>Anisotropy</i>	standard deviation of HOG features
			<i>Fourier Slope</i>	slope of line of best fit on the log-log plot of the Fourier power spectrum
Symmetry	(Brachmann and Redies 2016) (Bertamini and Rampone 2019)	2	<i>Fourier Sigma</i>	RMSE of line of best fit on the log-log plot of the Fourier power spectrum
			<i>Symmetry-LR</i>	left-right symmetry based on first layer activations on pre-trained AlexNet
CNN	(Brachmann, Barth, and Redies 2017)	2	<i>Symmetry-UD</i>	up-down symmetry based on first layer activations on pre-trained AlexNet
			<i>Sparseness</i>	median variance of each max-pooled response map from the first layer of pre-trained AlexNet
Other	(Savardi et al. 2018), (Zhang and Peng 2022)	>1	<i>Variability</i>	variance over all max-pooled response maps from the first layer of pre-trained AlexNet
			<i>Shot Scale</i>	close, medium, or long shot based on fine-tuned CNN
			<i>Bag of Objects</i>	frequency of objects detected by YOLOv5

Table 2: Eight groups of visual features extractable from thumbnails. The first seven are scalar; the last is discrete.

represents a labeled object, and the value corresponds to the frequency of that object in the image, analogous to a bag-of-words model in NLP.

## 4 Validations

Before further processing, we explored the ability of a feature to create consistent clusters of thumbnails.

### 4.1 Indoor/Outdoor.

Using the PlacesCNN model (Zhou et al. 2014), which can label each image as indoor/outdoor, we measured the entropy of each theme using the Gini index of statistical dispersion.

After cleaning the dataset, the Ukraine news story had 599 outdoor images and 575 indoor images with a Gini index of 0.50, a nearly perfect balance. Whereas, the COVID-19 news story had 931 indoor and 239 outdoor image with a Gini index of 0.33, mostly due to a strong predominance of medical press conferences.

For COVID-19, the Gini indices of two out of the three individual themes in Table 1 were less than the aggregate baseline, suggesting that our image clusters had produced groups that were more consistent in setting. This was especially true for the third “equipment” theme being quite pure (0.16) compared to the baseline (0.33). For Ukraine, the Gini indices also showed a decrease in entropy for all four individual themes in Table 1, particularly for the “rubble” theme (0.26), in which outdoor scenes dominated.

We considered these clustering results to be consistent.

### 4.2 CNN features.

We qualitatively evaluated the extraction of CNN features by visualizing images with the highest/lowest values for certain features.

Visual examination verified that images with low left-right symmetry tended to have a dominant subject framed at the edges of the screen; sparse images tended to be of low saturation; and images identified as close-up shots were usually portraits of people.

## 5 Results

### 5.1 Viewership behavior.

*Do U.S. and Chinese viewers behave differently on YouTube? We compared the distribution of viewership statistics (views, likes, comments) of each theme on each channel.*

**Views.** The log-log plots of views count vs. view rank strongly suggested a power law distribution. All four channels follow roughly the same behavior for the Ukraine event. For the COVID-19 news event, however, the Chinese channels show a noticeably steeper slope, suggesting a kind of Matthew effect, or possibly a media “contagion” effect due to the more immediate personal impact of this event there.

**Likes.** We did not find that likes adhered to a power law distribution. However, we did find that like-*rates* (likes divided by views) differed significantly between U.S. vs. Chinese channels; see Figure 5. U.S. channels generally have a lower like-rate, even though they have significantly more views overall.




Figure 4: Examples of top- and bottom-ranked thumbnails for three of the CNN-based features.

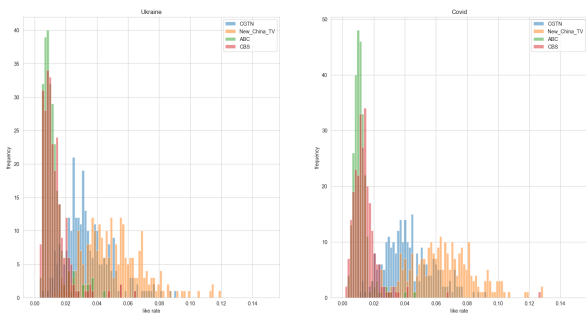


Figure 5: Like-rate distribution across events and channels.

**Comments.** We found no detectable differences in the rate at which viewers commented on a video.

## 5.2 Theme Preferences.

*Do U.S. and Chinese channels cover different news themes?*

**Ukraine.** As shown previously in Figure 3, U.S. thumbnails have a noticeably stronger focus on the two militaristic visual themes, whereas Chinese thumbnails focused more on the two more political visual themes.

This difference may reflect the intentions of the different media outlets, or the different audience preferences, or both. Chinese channels, due to political affiliation with Russia, might be more hesitant to cover the militaristic point of view, whereas sources based in the U.S., which have openly declared their support for Ukraine, may be more focused on bringing awareness to the human cost of the war. Alternatively, U.S. audiences may be more drawn to click on thumbnails that involve violence, and so these channels might cater to this preference to increase their metrics.

**COVID-19.** As shown previously in Figure 3, both cultures have a large emphasis on the medical context of COVID-19, with the “mask, equipment, syringe” visual theme being dominant.

**Themes over Time.** For both events, video uploads peak at times of crisis, and sometimes in a culture-specific way. For example, Figure 6 traces the COVID “mask” visual theme. The first and dominant peak occurs at the emergence of the Omicron variant (April, 2022), the second at the rescinding of the “Zero-COVID” policy in China (December, 2022) after which videos from the Chinese sources drop to zero, and the third at the Fall 2023 resurgence (September, 2023) reported only in the U.S.




Figure 6: Distribution of a COVID theme over time.

## 5.3 Aesthetic Preferences.

*Do U.S. vs. Chinese thumbnails differ in their aesthetics?*

**Feature Correlations.** We derived a feature correlation matrix between each image feature across the entire dataset, and note the following.

**Color** As expected, the HSV and the L\*a\*b\* spaces, which both attempt to match human visual perception, share several strong correlations, particularly with hue.

**HOG** The features of self-similarity and anisotropy are strongly negatively correlated. An image with low anisotropy (i.e., one with a lower standard deviation of edge intensities across orientations, due to many parallels within a building or to featureless regions within the sky) generates HOG features that are more likely to be similar across adjacent patches.

event	concept	hue_mean	saturation_mean	luminance_mean	lab(a)_mean	lab(b)_mean	color_entropy	aspect_ratio	image_size	contrast	luminance_entropy	self_similarity_neighbor	complexity	anisotropy	lr_symmetry	ud_symmetry	sparseness	variability	fourier_slope	fourier_sigma
ukraine	camouflage, army, soldier, gun, equipment	-0.11	0.06	-0.08	0.00	0.01	0.01	0.00	0.00	0.02	0.00	0.00	0.02	-0.04	-0.10	-0.08	0.06	-0.13	0.00	0.00
ukraine	speech, podium, give, deliver, microphone	-0.14	-0.02	-0.02	0.00	0.02	-0.01	-0.01	0.01	0.05	0.00	-0.02	-0.02	0.02	-0.08	-0.06	0.11	-0.16	0.04	0.00
ukraine	business suit, news, interview, woman, suit	<b>-0.07</b>	<b>0.16</b>	<b>-0.11</b>	<b>0.01</b>	<b>0.01</b>	-0.01	0.00	0.00	0.00	0.00	<b>-0.04</b>	<b>-0.08</b>	<b>0.03</b>	<b>-0.17</b>	<b>-0.07</b>	<b>0.14</b>	-0.10	<b>0.03</b>	<b>0.00</b>
ukraine	debris, city, rubble, destruction, damage	-0.01	0.06	0.00	0.00	0.00	0.00	0.01	0.01	0.01	0.00	<b>-0.02</b>	-0.04	0.01	<b>-0.10</b>	<b>-0.05</b>	0.05	-0.16	<b>0.06</b>	<b>0.00</b>
covid	interview, sit, table, blue, speech	<b>-0.09</b>	-0.04	<b>-0.13</b>	0.00	<b>0.04</b>	<b>0.04</b>	0.00	0.00	<b>0.06</b>	0.00	0.01	0.02	-0.02	<b>-0.18</b>	<b>-0.09</b>	<b>0.09</b>	-0.02	<b>0.03</b>	<b>0.00</b>
covid	mask, equipment, syringe, hand, liquid	<b>-0.06</b>	<b>0.04</b>	<b>-0.04</b>	<b>0.01</b>	<b>0.00</b>	0.00	<b>0.01</b>	<b>0.00</b>	<b>0.03</b>	0.00	<b>-0.03</b>	<b>-0.06</b>	<b>0.04</b>	<b>-0.10</b>	<b>-0.05</b>	<b>0.11</b>	<b>-0.07</b>	<b>0.05</b>	<b>0.00</b>
covid	business suit, tie, suit, interview, microphone	0.02	<b>0.08</b>	<b>-0.11</b>	0.00	0.00	0.01	0.00	0.00	0.01	<b>0.01</b>	-0.01	0.03	-0.03	<b>-0.13</b>	<b>-0.13</b>	<b>0.11</b>	-0.02	-0.01	0.00

Figure 7: Preferences for image aesthetics. Blue: U.S. is greater. Red: Chinese is greater. Bold: statistically significant at  $p < .05$ .

*Dimension vs. Fourier* An inverse relationship between image size and Fourier sigma was unexpected, and is likely an artifact of pre-processing. Because we had to remove user-added black bars of various sizes, thumbnails now have different dimensions, which then affects their spatial frequencies.

**Aesthetic Comparisons, Scalar.** To examine cultural preferences, we display in Figure 7 the results from a two-sample (U.S. vs. Chinese) t-test for the 19 scalar visual features, for each of the seven visual themes given in Table 1.

The value in each cell gives the normalized difference (“difference over sum”) between the U.S. and Chinese average feature values for that cell. Positive values indicate that U.S. thumbnails exhibit the larger average value for the feature, and their cell is shaded blue; negative values in red cells indicate the reverse. Values are made bold when the difference between U.S. and Chinese thumbnail distributions is statistically significant ( $p < 0.05$ ).

We note that some features show a cultural preference that generalizes across visual themes, irrespective of the image content. These show up in Figure 7 as (nearly) uniformly colored columns.

*Color: Saturation.* Although U.S. images are less “colorful”, their colors tend to be more saturated. This is possibly the result of an editing process, to give the impression that they had been photographed and edited separately, whereas Chinese thumbnails may strive to be more candid, as if they were taken directly from a frame in the video.

*Lightness: Luminance.* In general, Chinese thumbnails also tend to be brighter, and appear less formal overall, occasionally even overlaying bright animated components on top of background images.

*Fourier: Fourier Slope.* This metric is greater in U.S. thumbnails, especially in ones relating to “rubble” and “debris” in the Ukraine news story. The increased detail may be a result of higher-grade cameras in the field, further supporting an observation that U.S. thumbnails are more professional, whereas Chinese ones are less formal.

*Symmetry: -LR, -UD.* Chinese thumbnails are, on average, more symmetric in both the horizontal and vertical directions, across all events, and statistically significantly so in five of the seven visual themes. We note that U.S. thumbnails instead adhere more closely to the photographic “rule of thirds” where the foreground figure is framed more to the left or right. Chinese thumbnails appear to come from the actual video footage itself, which are more candid in their approach, and favor long balanced landscape shots.

*CNN: Sparseness, Variability.* U.S. thumbnails are more sparse, across all seven visual themes. We note that many U.S. thumbnails are portraits taken with cameras set at a low aperture, causing the (usually dark) background to be blurred; blurred regions have low variance. This may also explain why Chinese thumbnails score higher in the variability metric.

**Aesthetic Comparisons, Discrete.** Examples of the two “Other” features follow.

*Other: Shot Scale.* U.S. thumbnails have a stronger preference for close-up shots, portraits in most cases, whereas Chinese thumbnails favor long shots. See Figure 9 for their comparative use in the four visual themes in the Ukraine news story. This may suggest that U.S. channels place more importance on specific individuals making news.

*Other: Bag of Objects.* A comparison of the objects that appear most often in thumbnails can be displayed with a word cloud, as in Figure 10, for events vs. cultures.

All four are dominated by common human-associated objects, such as a person, tie, chair, laptop, TV, bottle, cup, car,

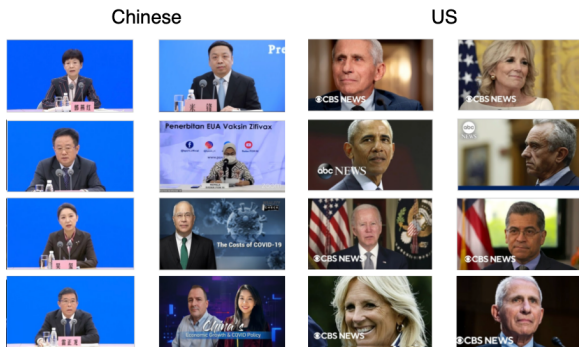


Figure 8: Chinese vs. U.S. press conferences; note blue.

*Color: Hue.* Chinese thumbnails tend to have greater hue values (i.e., they are “more colorful”) than U.S. thumbnails, especially for thumbnails depicting interviews or press conferences, which favor featureless but vibrant blue backdrops; see Figure 8. U.S. thumbnails for the same visual theme often instead favor more natural backdrops such as a room or the White House.





## References

- Anwar, A.; Kanwal, S.; Tahir, M.; Saqib, M.; Uzair, M.; Rahmani, M. K. I.; and Ullah, H. 2021. A survey on image aesthetic assessment. *arXiv preprint arXiv:2103.11616*.
- Aydn, T. O.; Smolic, A.; and Gross, M. 2014. Automated aesthetic analysis of photographic images. *IEEE transactions on visualization and computer graphics*, 21(1): 31–42.
- Bartho, R.; Thoemmes, K.; and Redies, C. 2023. Predicting beauty, liking, and aesthetic quality: A comparative analysis of image databases for visual aesthetics research. *arXiv preprint arXiv:2307.00984*.
- Bertamini, M.; and Rampone, G. 2019. *The study of symmetry in empirical aesthetics*, chapter 22. The Oxford Handbook of Empirical Aesthetics. Oxford University Press.
- Bhargava, R.; Bishop, C.; and Zuckerman, E. 2020. Mapping and Visualizing News Images for Media Research. In *Proceedings of Computation+ Journalism Symposium*.
- Bosch, A.; Zisserman, A.; and Munoz, X. 2007. Representing shape with a spatial pyramid kernel. In *Proceedings of the 6th ACM international conference on Image and video retrieval*, 401–408.
- Brachmann, A.; Barth, E.; and Redies, C. 2017. Using CNN features to better understand what makes visual artworks special. *Frontiers in psychology*, 8: 266181.
- Brachmann, A.; and Redies, C. 2016. Using convolutional neural network filters to measure left-right mirror symmetry in images. *Symmetry*, 8(12): 144.
- Braun, J.; Amirshahi, S. A.; Denzler, J.; and Redies, C. 2013. Statistical image properties of print advertisements, visual artworks and images of architecture. *Frontiers in Psychology*, 4: 60395.
- Burgess, J.; and Green, J. 2018. *YouTube: Online video and participatory culture*. John Wiley & Sons.
- Cao, J.; Qi, P.; Sheng, Q.; Yang, T.; Guo, J.; and Li, J. 2020. Exploring the role of visual content in fake news detection. *Disinformation, Misinformation, and Fake News in Social Media: Emerging Research Challenges and Opportunities*, 141–161.
- Chen, W.; Rabhi, F.; Liao, W.; and Al-Qudah, I. 2023. Leveraging state-of-the-art topic modeling for news impact analysis on financial markets: A comparative study. *Electronics*, 12(12): 2605.
- Datta, R.; Joshi, D.; Li, J.; and Wang, J. Z. 2006. Studying aesthetics in photographic images using a computational approach. In *Computer Vision—ECCV 2006: 9th European Conference on Computer Vision, Graz, Austria, May 7–13, 2006, Proceedings, Part III 9*, 288–301. Springer.
- Doshi, N.; Shikkenawis, G.; and Mitra, S. K. 2020. Image aesthetics assessment using multi channel convolutional neural networks. In *Computer Vision and Image Processing: 4th International Conference, CVIP 2019, Jaipur, India, September 27–29, 2019, Revised Selected Papers, Part II 4*, 15–24. Springer.
- Fadaei, S. 2021. Comparison of color spaces in dcd-based content-based image retrieval systems. In *2021 7th International Conference on Signal Processing and Intelligent Systems (ICSPIS)*, 1–4. IEEE.
- Fahmy, S. 2010. Contrasting visual frames of our times: A framing analysis of English-and Arabic-language press coverage of war and terrorism. *International communication gazette*, 72(8): 695–717.
- FORCE11. 2020. The FAIR Data principles. <https://force11.org/info/the-fair-data-principles>. Accessed: 2024-03-24.
- Gebbru, T.; Morgenstern, J.; Vecchione, B.; Vaughan, J. W.; Wallach, H.; Iii, H. D.; and Crawford, K. 2021. Datasheets for datasets. *Communications of the ACM*, 64(12): 86–92.
- Geller, H. A.; Bartho, R.; Thömmes, K.; and Redies, C. 2022. Statistical image properties predict aesthetic ratings in abstract paintings created by neural style transfer. *Frontiers in neuroscience*, 16: 999720.
- Grootendorst, M. P. 2021. Concept. [maartengr.github.io/Concept/index.html](https://maartengr.github.io/Concept/index.html). Accessed: 2024-01-04.
- Hameleers, M.; Powell, T. E.; Van Der Meer, T. G.; and Bos, L. 2020. A picture paints a thousand lies? The effects and mechanisms of multimodal disinformation and rebuttals disseminated via social media. *Political communication*, 37(2): 281–301.
- Huang, X.; Zhang, Y.; Ma, J.; Tian, W.; Feng, R.; Zhang, Y.; Li, Y.; Guo, Y.; and Zhang, L. 2023. Tag2text: Guiding vision-language model via image tagging. *arXiv preprint arXiv:2303.05657*.
- Igaya, K.; Yi, S.; Wahle, I. A.; Tanwisuth, K.; and O’Doherty, J. P. 2021. Aesthetic preference for art can be predicted from a mixture of low-and high-level visual features. *Nature Human Behaviour*, 5(6): 743–755.
- Joffe, H. 2008. The power of visual material: Persuasion, emotion and identification. *Diogenes*, 55(1): 84–93.
- Kao, Y.; He, R.; and Huang, K. 2017. Deep aesthetic quality assessment with semantic information. *IEEE Transactions on Image Processing*, 26(3): 1482–1495.
- Kim, H.; Kim, J.; Han, Y.; Hong, H.; Kwon, O.-S.; Park, Y.-W.; Elmqvist, N.; Ko, S.; and Kwon, B. C. 2023. Towards Visualization Thumbnail Designs That Entice Reading Data-Driven Articles. *IEEE Transactions on Visualization and Computer Graphics*.
- Koh, B.; and Cui, F. 2022. An exploration of the relation between the visual attributes of thumbnails and the view-through of videos: The case of branded video content. *Decision Support Systems*, 160: 113820.
- Li, C.; and Chen, T. 2009. Aesthetic visual quality assessment of paintings. *IEEE Journal of selected topics in Signal Processing*, 3(2): 236–252.
- Lo, K.-Y.; Liu, K.-H.; and Chen, C.-S. 2013. Intelligent photographing interface with on-device aesthetic quality assessment. In *Computer Vision-ACCV 2012 Workshops: ACCV 2012 International Workshops, Daejeon, Korea, November 5–6, 2012, Revised Selected Papers, Part II 11*, 533–544. Springer.
- Luo, W.; Wang, X.; and Tang, X. 2011. Content-based photo quality assessment. In *2011 international conference on computer vision*, 2206–2213. IEEE.

- Machajdik, J.; and Hanbury, A. 2010. Affective image classification using features inspired by psychology and art theory. In *Proceedings of the 18th ACM international conference on Multimedia*, 83–92.
- Mallon, B.; Redies, C.; and Hayn-Leichsenring, G. U. 2014. Beauty in abstract paintings: perceptual contrast and statistical properties. *Frontiers in human neuroscience*, 8: 161.
- Martikainen, J.; and Sakki, I. 2021. How newspaper images position different groups of people in relation to the COVID-19 pandemic: A social representations approach. *Journal of Community & Applied Social Psychology*, 31(4): 465–494.
- Mather, G. 2018. Visual image statistics in the history of Western art. *Art & Perception*, 6(2-3): 97–115.
- Murray, N.; Marchesotti, L.; and Perronnin, F. 2012. AVA: A large-scale database for aesthetic visual analysis. In *2012 IEEE conference on computer vision and pattern recognition*, 2408–2415. IEEE.
- Peli, E. 1990. Contrast in complex images. *JOSA A*, 7(10): 2032–2040.
- Peng, Y.; and JEMMOTT III, J. B. 2018. Feast for the Eyes: Effects of Food Perceptions and Computer Vision Features on Food Photo Popularity. *International Journal of Communication (19328036)*, 12.
- Rafiee, A.; Spooren, W.; and Sanders, J. 2023. Framing similar issues differently: a cross-cultural discourse analysis of news images. *Social Semiotics*, 33(3): 515–538.
- Redi, M.; Rasiwasia, N.; Aggarwal, G.; and Jaimes, A. 2015. The beauty of capturing faces: Rating the quality of digital portraits. In *2015 11th IEEE international conference and workshops on automatic face and gesture recognition (FG)*, volume 1, 1–8. IEEE.
- Redies, C.; Hasenstein, J.; Denzler, J.; et al. 2008. Fractal-like image statistics in visual art: similarity to natural scenes. *Spatial vision*, 21(1-2): 137–148.
- Savardi, M.; Signoroni, A.; Migliorati, P.; and Benini, S. 2018. Shot scale analysis in movies by convolutional neural networks. In *2018 25th IEEE International Conference on Image Processing (ICIP)*, 2620–2624. IEEE.
- Schifanella, R.; Redi, M.; and Aiello, L. M. 2015. An image is worth more than a thousand favorites: Surfacing the hidden beauty of flickr pictures. In *Proceedings of the international AAAI conference on web and social media*, 397–406.
- Segalin, C.; Perina, A.; Cristani, M.; and Vinciarelli, A. 2016. The pictures we like are our image: continuous mapping of favorite pictures into self-assessed and attributed personality traits. *IEEE Transactions on Affective Computing*, 8(2): 268–285.
- Sidhu, D. M.; McDougall, K. H.; Jalava, S. T.; and Bodner, G. E. 2018. Prediction of beauty and liking ratings for abstract and representational paintings using subjective and objective measures. *PLoS one*, 13(7): e0200431.
- Smith, N.; and Joffe, H. 2013. How the public engages with global warming: A social representations approach. *Public Understanding of Science*, 22(1): 16–32.
- Thieleking, R.; Medawar, E.; Disch, L.; and Witte, A. V. 2020. Art. pics database: an open access database for art stimuli for experimental research. *Frontiers in Psychology*, 11: 576580.
- Thomson, T.; Perreault, G.; and Duffy, M. 2018. Politicians, photographers, and a Pope: How state-controlled and independent media covered Francis’s 2015 Cuba visit. *Journalism Studies*, 19(9): 1313–1330.
- Wallbing, S. 2021. Computational Analysis of Swedish Newspapers Using Topic Detection and Sentiment Analysis.
- Xie, L.; Natsev, A.; Kender, J. R.; Hill, M.; and Smith, J. R. 2011. Visual memes in social media: tracking real-world news in youtube videos. In *Proceedings of the 19th ACM international conference on Multimedia*, 53–62.
- Zhang, H.; and Peng, Y. 2022. Image clustering: An unsupervised approach to categorize visual data in social science research. *Sociological Methods & Research*, 00491241221082603.
- Zhou, B.; Lapedriza, A.; Xiao, J.; Torralba, A.; and Oliva, A. 2014. Learning deep features for scene recognition using places database. *Advances in neural information processing systems*, 27.

## Ethics Checklist

1. For most authors...
  - (a) Would answering this research question advance science without violating social contracts, such as violating privacy norms, perpetuating unfair profiling, exacerbating the socio-economic divide, or implying disrespect to societies or cultures? **Yes**
  - (b) Do your main claims in the abstract and introduction accurately reflect the paper's contributions and scope? **Yes**
  - (c) Do you clarify how the proposed methodological approach is appropriate for the claims made? **Yes**
  - (d) Do you clarify what are possible artifacts in the data used, given population-specific distributions? **Yes**
  - (e) Did you describe the limitations of your work? **Yes**
  - (f) Did you discuss any potential negative societal impacts of your work? **No**. The work is purely observational.
  - (g) Did you discuss any potential misuse of your work? **No**. No system was designed or created.
  - (h) Did you describe steps taken to prevent or mitigate potential negative outcomes of the research, such as data and model documentation, data anonymization, responsible release, access control, and the reproducibility of findings? **No**. Data is publicly available from news channels and has not identified individuals without their consent.
  - (i) Have you read the ethics review guidelines and ensured that your paper conforms to them? **Yes**
2. Additionally, if your study involves hypotheses testing...
  - (a) Did you clearly state the assumptions underlying all theoretical results? **No** theoretical results.
  - (b) Have you provided justifications for all theoretical results? **No** theoretical results.
  - (c) Did you discuss competing hypotheses or theories that might challenge or complement your theoretical results? **No** theoretical results.
  - (d) Have you considered alternative mechanisms or explanations that might account for the same outcomes observed in your study? **Yes**
  - (e) Did you address potential biases or limitations in your theoretical framework? **Yes**. Dataset is limited in size.
  - (f) Have you related your theoretical results to the existing literature in social science? **No** theoretical results.
  - (g) Did you discuss the implications of your theoretical results for policy, practice, or further research in the social science domain? **No** theoretical results.
3. Additionally, if you are including theoretical proofs...
  - (a) Did you state the full set of assumptions of all theoretical results? **No** theoretical proofs.
  - (b) Did you include complete proofs of all theoretical results? **No** theoretical proofs.
4. Additionally, if you ran machine learning experiments...
  - (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? **No**. No substantive code was created, and the off-the-shelf CNN is cited. Data is publicly available via YouTube API.
  - (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? **Yes**
  - (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? **No**. Experiments were statistical and run once.
  - (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? **No**. No training of any system was involved.
  - (e) Do you justify how the proposed evaluation is sufficient and appropriate to the claims made? **Yes**.
  - (f) Do you discuss what is "the cost" of misclassification and fault (in)tolerance? **No**. The paper notes preferences and no hard decisions.
5. Additionally, if you are using existing assets (e.g., code, data, models) or curating/releasing new assets, **without compromising anonymity**...
  - (a) If your work uses existing assets, did you cite the creators? **Yes**
  - (b) Did you mention the license of the assets? **No**. The system (Cluster) and data source (YouTube) are common and well-known.
  - (c) Did you include any new assets in the supplemental material or as a URL? **No** new assets.
  - (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? **No** consent was required.
  - (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? **No**. Data was standard YouTube download off well-established news channels.
  - (f) If you are curating or releasing new datasets, did you discuss how you intend to make your datasets FAIR (see FORCE11 (2020))? **No** new datasets.
  - (g) If you are curating or releasing new datasets, did you create a Datasheet for the Dataset (see Gebru et al. (2021))? **No** new datasets.
6. Additionally, if you used crowdsourcing or conducted research with human subjects, **without compromising anonymity**...
  - (a) Did you include the full text of instructions given to participants and screenshots? **No** human subjects.
  - (b) Did you describe any potential participant risks, with mentions of Institutional Review Board (IRB) approvals? **No** human subjects.
  - (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? **No** human subjects.
  - (d) Did you discuss how data is stored, shared, and de-identified? **No** human subjects.