

Technical Disclosure Commons

Defensive Publications Series

November 2022

Automatic Adaptive Image Cropping to Preserve Regions of Interest

Zhuowei Zhang

William Leszczuk

Jerome Miglino

Chris Ye

Follow this and additional works at: https://www.tdcommons.org/dpubs_series

Recommended Citation

Zhang, Zhuowei; Leszczuk, William; Miglino, Jerome; and Ye, Chris, "Automatic Adaptive Image Cropping to Preserve Regions of Interest", Technical Disclosure Commons, (November 04, 2022)
https://www.tdcommons.org/dpubs_series/5467



This work is licensed under a [Creative Commons Attribution 4.0 License](https://creativecommons.org/licenses/by/4.0/).

This Article is brought to you for free and open access by Technical Disclosure Commons. It has been accepted for inclusion in Defensive Publications Series by an authorized administrator of Technical Disclosure Commons.

Automatic Adaptive Image Cropping to Preserve Regions of Interest

ABSTRACT

This disclosure describes techniques to adaptively crop an image to fit fixed display areas while preserving regions of interest of the image. A bounding box of a region of interest is obtained for the image and a sorted list of candidate aspect ratios for the image are determined. Candidate image sizes are determined for candidate aspect ratios and are compared to the bounding box size. If a candidate image has a greater size in a crop axis than the bounding box, then that aspect ratio is selected for use in the crop operation. This ensures that the region of interest is preserved in the cropped image.

KEYWORDS

- Image crop
- Aspect ratio
- Adaptive crop
- Letterboxing
- Thumbnail
- Image preview

BACKGROUND

Images may be cropped so that they fit within particular display areas having aspect ratios that are different from that of the images. For example, images may be cropped when displayed (e.g., as thumbnail images) in menus, search results, image grids, albums, web pages, and other presentations in which the images are displayed in fixed areas so that they align with each other and/or with margins, text, displayed objects, etc.

A source image may be cropped to a given aspect ratio of the display area. The cropped image may be zoomed and/or letterboxed (e.g., with black bars displayed on either or both sides of the image) to fit in the display area. However, a region of interest of the cropped image (e.g., main subject depicted in the image) may be altered by the crop operation. For example, the region of interest may have portions that are least partially cut off at the edges of the display area in order to fit the image into the display area.

DESCRIPTION

This disclosure describes techniques for adaptive cropping of an image based on a region of interest in the image. A size of the region of interest in the image is compared to image sizes determined from a list of candidate aspect ratios. A candidate aspect ratio that allows the cropped image to preserve the region of interest and fit in a desired display area is selected. Described techniques can determine which of several candidate aspect ratios to use for the crop operation. These features allow an image to be cropped to fit in a desired display area while preserving all of the region of interest in the original image. For example, images cropped using the described techniques can be provided as thumbnails in search results, menus, grids, or other presentations that include portrait and/or landscape images in fixed size display areas and present a unified appearance of the images.

The described techniques can be implemented on any suitable device, e.g., desktop or laptop computer, portable user device (e.g., a smartphone), server device, etc. The user is provided with options to enable or disable described techniques. The user can permit specific data or types of data (e.g., specific images) to be processed and can deny processing of other data or types of data.

Determining candidate aspect ratios

According to described features, a source image that is to be displayed in a particular fixed-size display area is obtained. The width and height of the source image is determined. For example, a user can input the source image, the source image can be obtained from search results, the source image can be a photo captured by a camera of a user device or by a different camera, etc.

A region of interest (ROI) of the source image is identified. The region of interest is the main subject of the source image. It is defined by a rectangular bounding box. The width and height of the bounding box is determined. In some examples, the region of interest of an image showing a person is the face of the person. An image showing a group of people may have a region of interest that includes all the faces of the people, or the region may include only the largest faces of people in the image, e.g., if other faces are small (e.g., further distance away from the camera that captured the image). An image of a landscape may have a region of interest that covers most or all of the image, e.g., if there are not any prominent objects depicted in the image. The region of interest and bounding box can be obtained using any of various techniques. Such techniques can include, for example, image feature detection techniques, face detection techniques, other image recognition techniques, a saliency map of the image, etc.

The aspect ratio of the display area and the aspect ratio of the source image are determined, e.g., as a ratio of width to height. A set of candidate aspect ratios is then determined. For example, the candidate aspect ratios in the set can be specified or predetermined, such as commonly used aspect ratios (e.g., 1:1, 2:3, 4:3, 16:9, etc.). In addition, the candidate aspect ratios also include the aspect ratios of the display area and of the source image. Optionally, a filter can be used to filter out the widest and/or narrowest aspect ratios allowed in the set of

candidate aspect ratios. For example, if aspect ratios narrower than a square (1:1 aspect ratio) are not desired, such aspect ratios are not included in the candidate set.

It is then determined which axis of the source image will fill the corresponding axis of the display area (the fill axis), and which axis of the source image will be cropped (the crop axis). The cropped image can be zoomed, if necessary, to fill this axis of the display area. The crop axis is the dimension of the source image that is to be cropped to fit into the display area.

The fill axis can be determined by comparing the source image aspect ratio and the display area aspect ratio. If the source image is wider than the display area, e.g., the source image aspect ratio is greater than the display area aspect ratio, then the horizontal (width) axis of the source image is the crop axis and the vertical axis is the fill axis. In this case, the set of candidate aspect ratios is sorted from smallest (narrowest) to largest (widest).

Otherwise, e.g., if the source image aspect ratio is lesser than the display area aspect ratio such that the width of the source image is narrower than the display area, then the vertical (height) axis is the crop axis and the horizontal axis is the fill axis. In this case, the set of candidate aspect ratios is sorted from largest (widest) to smallest (narrowest). Other sorting orders can alternatively be used, e.g., a particular order in which preferred aspect ratios appear earlier in the list.

Finding the candidate aspect ratio that preserves the ROI bounding box

After the set of candidate aspect ratios has been sorted, candidate aspect ratios are selected in the sort order and tested until a candidate aspect ratio satisfies the test, where the satisfying aspect ratio is selected for use in the crop operation.

Each selected candidate aspect ratio is tested by first computing a size of a candidate image that has the selected candidate aspect ratio. The fill axis of the candidate image has a size

that is 100% of the size of the fill axis of the source image. The crop axis of the candidate image has a size that can be determined as a percentage of the crop axis of the source image. If the crop axis is horizontal, its size is determined as the candidate aspect ratio divided by the source image aspect ratio, which can be multiplied by 100% to obtain a percentage. If the crop axis is vertical, its size is determined as the source image aspect ratio divided by the candidate aspect ratio, multiplied by 100%.

The size of the ROI bounding box along the axis that corresponds to the crop axis is converted to a percentage of the source image size along the crop axis. The bounding box size percentage is then compared to the candidate image size percentage along the crop axis. If the size of the candidate image along the crop axis is equal to the bounding box size, or greater than the bounding box size within a threshold amount (or a threshold multiplier), then the candidate aspect ratio of the candidate image is selected as the resulting aspect ratio to be used in cropping the source image to fit in the display area. This aspect ratio is the first aspect ratio in the list of candidate aspect ratios that preserves all of the bounding box area of the source image.


If the size of the candidate image along the crop axis is less than the bounding box size or is greater than the bounding box size by more than the threshold amount, then the next candidate aspect ratio in the sorted list of candidate aspect ratios is selected and is tested similarly as described above. If necessary, this is repeated for further candidate aspect ratios in the list until a candidate aspect ratio is found that satisfies the test and is used as the resulting aspect ratio.

The source image is then cropped according to the resulting aspect ratio. For example, the crop operation can be centered on the center of the bounding box. The cropped image is zoomed and/or letterboxed as needed so that the fill axis of the cropped image entirely fills the corresponding axis of the display area.

The threshold used in the comparison described above can be configurable to allow different amounts of area outside the ROI bounding box (e.g., a margin) to be present in the cropped image along the crop axis. For example, if a first threshold is greater than a second threshold, the first threshold allows a greater amount of margin area outside the bounding box in the cropped image to satisfy the test.

Examples


Examples of the above techniques are presented below. In a first example, a portrait source image is cropped to fit into a display area that has a landscape aspect ratio. In a second example, a landscape source image is cropped to fit into a display area having a portrait aspect ratio.



Figs. 1 and 2: Example source image with bounding box and target display area

Fig. 1 shows a source image (101) for which a bounding box (102) has been determined based on the content of the source image. In this example, the bounding box surrounds the face of the depicted person, where the face has been detected as the region of interest for this image. The source image is a portrait image that has a greater height than width.

Fig. 2 shows that the source image is to be cropped to fit into a display area (201) that has a landscape aspect ratio with a width greater than its height. In this example, the display area aspect ratio is 16:9. The display area is shown superimposed on the source image to indicate that upper and lower portions of the bounding box will be cut off if the source image is cropped to have the width of the display area.



Figs. 3 and 4: Selected aspect ratio and cropped image in display area

For Fig. 3, a number of candidate aspect ratios have been sorted into a list from widest to narrowest and have been tested as described above, where the crop axis is the vertical axis. If a candidate aspect ratio is wider than the display area or is narrower than the source image, it is skipped. Otherwise, a candidate image is determined based on the candidate aspect ratio and its height is compared to the bounding box height. If the candidate image height is equal to or greater than the bounding box height within a threshold, it is selected as the resulting aspect ratio. Fig. 3 shows a candidate image size (301) that has an aspect ratio of 4:3, which has a height that is greater than the height of the bounding box by an amount within the threshold; thus, this aspect ratio is selected for use in the crop operation.

Fig. 4 shows the cropped image (401) that has been cropped from the source image based on the selected 4:3 aspect ratio and is displayed within the display area. The cropped image is letterboxed on the left and right sides (as indicated by the black bands) so that it is centered within the 16:9 display area.




Fig. 5: Example source image with bounding box




Fig. 6: Source image cropped to fit display area

In the second example, as shown in Fig. 5, a source image (501) can include a bounding box (502) that has been determined based on the content of the source image. The bounding box surrounds the face of the depicted person similarly to Fig. 1. The source image is a landscape image that has a greater width than height.

Fig. 6 shows that the source image is to be cropped to fit into a display area (601) that has a portrait aspect ratio with a height greater than its width. In this example, the portrait area aspect ratio is 9:16. The display area is shown superimposed on the source image to indicate that left and right portions of the bounding box will be cut off if the source image is cropped to have the height of the display area.




Fig. 7: Image with selected aspect ratio




Fig. 8: Cropped image in display area

For Fig. 7, a number of candidate aspect ratios have been sorted into a list from narrowest to widest and have been tested as described above, where the crop axis is the horizontal axis. If a candidate aspect ratio is narrower than the display area or is wider than the source image, it is skipped. Otherwise, a candidate image is determined based on the candidate aspect ratio and its width is compared to the bounding box width. If the candidate image width is equal to or greater than the bounding box width within a threshold, it is selected as the resulting aspect ratio. Fig. 7 shows a candidate image size (701) that has an aspect ratio of 3:4, which has a height that is greater than the height of the bounding box by an amount within the threshold; thus, this aspect ratio is selected for use in the crop operation.

Fig. 8 shows the cropped image (801) that has been cropped from the source image based on the selected 3:4 aspect ratio and is displayed within the display area. The cropped image is letterboxed at the top and bottom so that it is centered within the 9:16 display area.

Users are provided with options to grant permissions to and/or to disable described features entirely. The various features of the system are implemented only with user permission to access user information that serves as input to the system (e.g., user images, user context information, camera input, user's preferences, etc.). Users may be provided with controls

allowing the user to make an election as to both if and when systems, programs or features described herein may enable collection of user information, and if the user is sent content or communications from a server. Certain techniques are not implemented if users deny permission. In addition, certain data may be treated in one or more ways before it is stored or used, so that personally identifiable information is removed. For example, a user's identity may be treated so that no personally identifiable information can be determined for the user. Thus, the user may have control over what information is collected about the user, how that information is used, and what information is provided to the user.

CONCLUSION

This disclosure describes techniques to adaptively crop an image to fit fixed display areas while preserving regions of interest of the image. A bounding box of a region of interest is obtained for the image and a sorted list of candidate aspect ratios for the image are determined. Candidate image sizes are determined for candidate aspect ratios and are compared to the bounding box size. If a candidate image has a greater size in a crop axis than the bounding box, then that aspect ratio is selected for use in the crop operation. This ensures that the region of interest is preserved in the cropped image.