



Figure 6.1: (Left) Naturally colored image of tiger in grass; (right) with transformed colors, recognition of a tiger is less secure – perhaps it's a cat on a rug?



Figure 6.9: (Left) Input RGB image; (center) saturation S increased by 40%; (right) saturation S decreased by 20%. (Photo by Frank Biocca.)



Figure 6.10: “White pixels” are segmented from the color image at the left. Individual connected components of white pixels are arbitrarily labeled by a coloring algorithm as described in Chapter 3. (Analysis contributed by David Moore.)

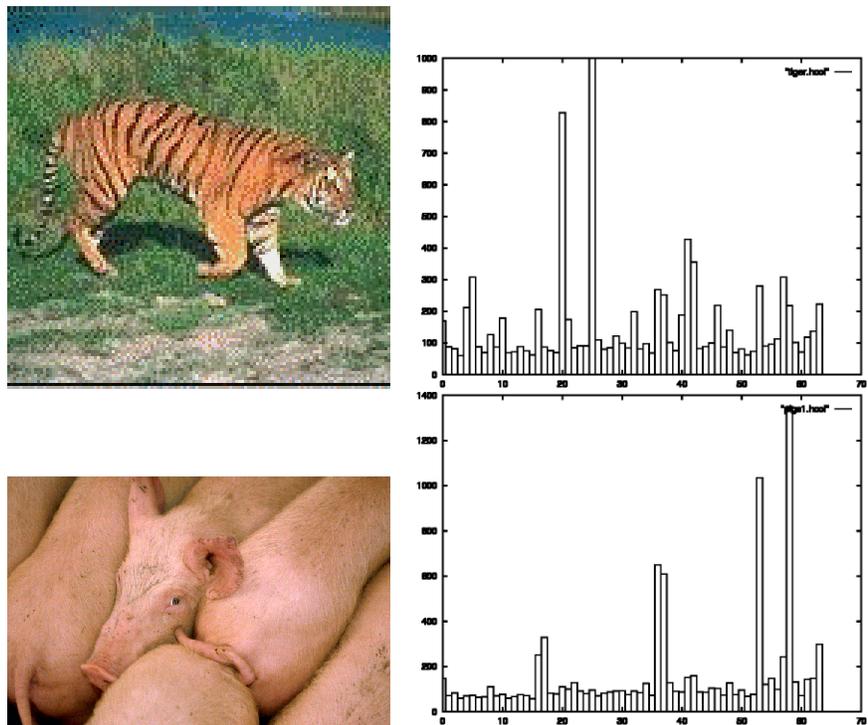


Figure 6.11: Color images and their 64-bin histograms (obtained from A. Vailaya).

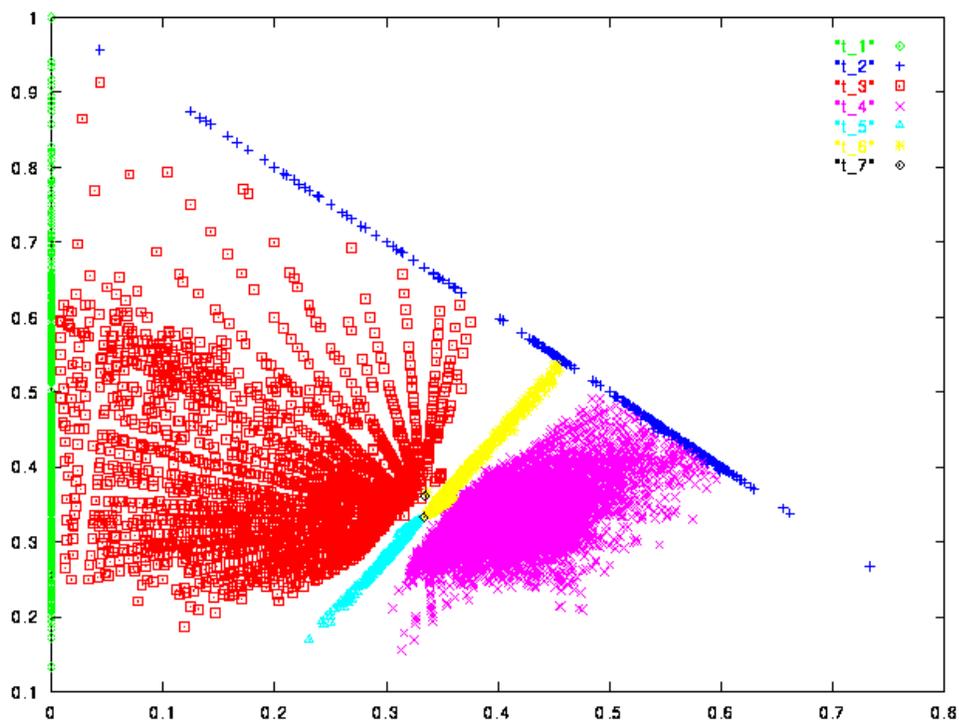


Figure 6.12: Skin color clusters obtained from training: the horizontal axis is R_{norm} and the vertical axis is G_{norm} . The cluster labeled as τ_4 is the primary face color, clusters τ_5 and τ_6 are secondary face clusters associated with shadowed or bearded areas of a face. (Figure from V. Bakic.)

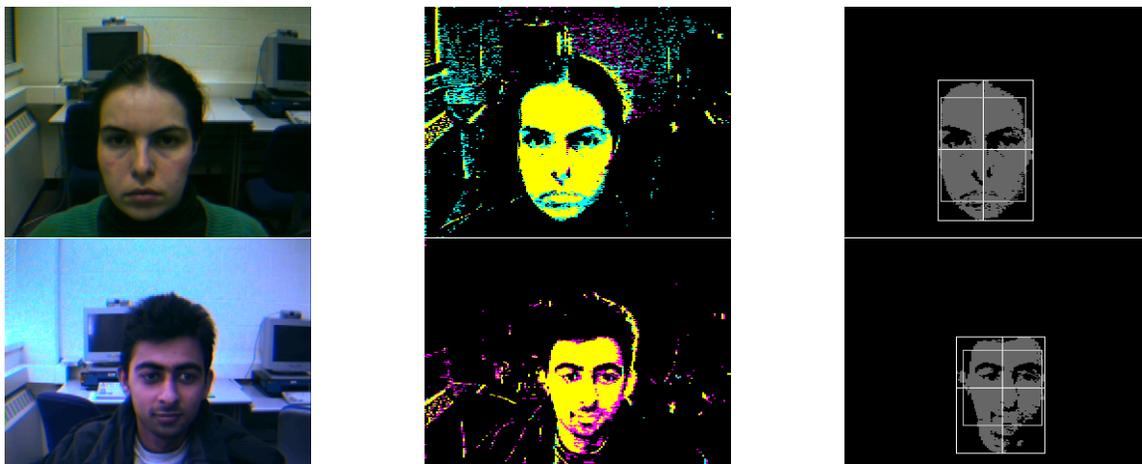


Figure 6.13: Face extraction examples: (left) input image, (middle) labeled image, (right) boundaries of the extracted face region. (Images from V. Bakic.)

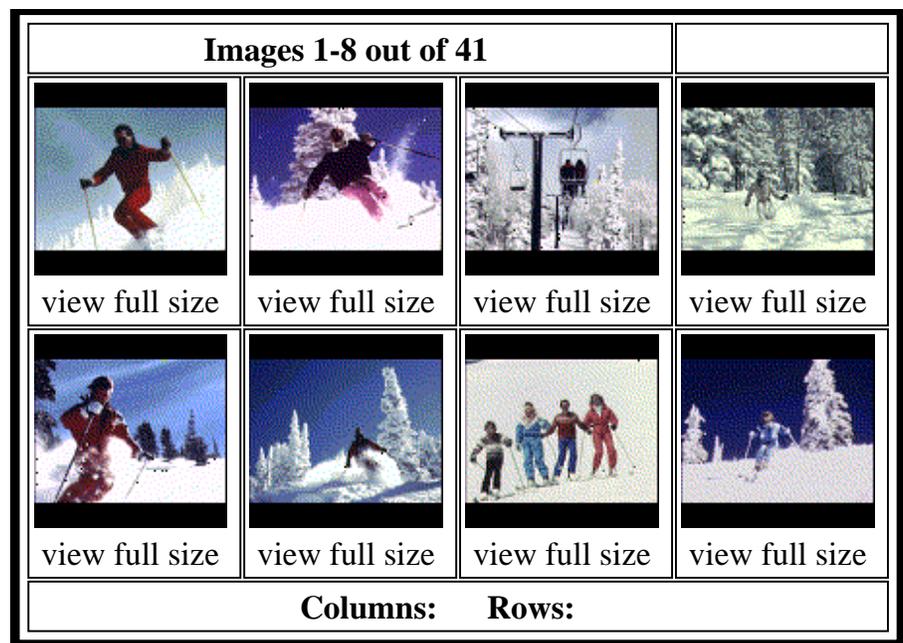


Figure 8.3: Results of a QBIC search based on color layout similarity; the query is the example image shown in the top left position.

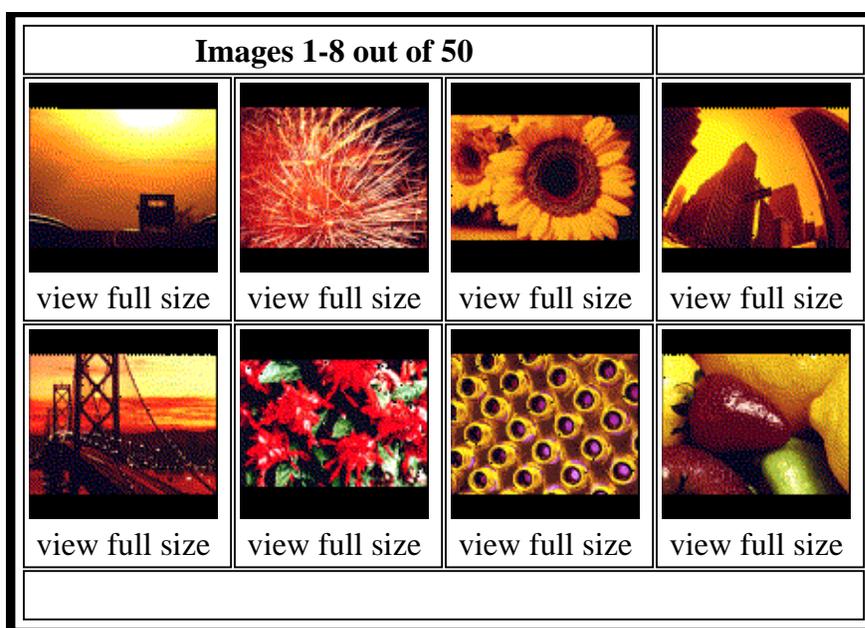


Figure 8.4: Results of a QBIC search based on color percentages; the query specified 40% red, 30% yellow, and 10% black.

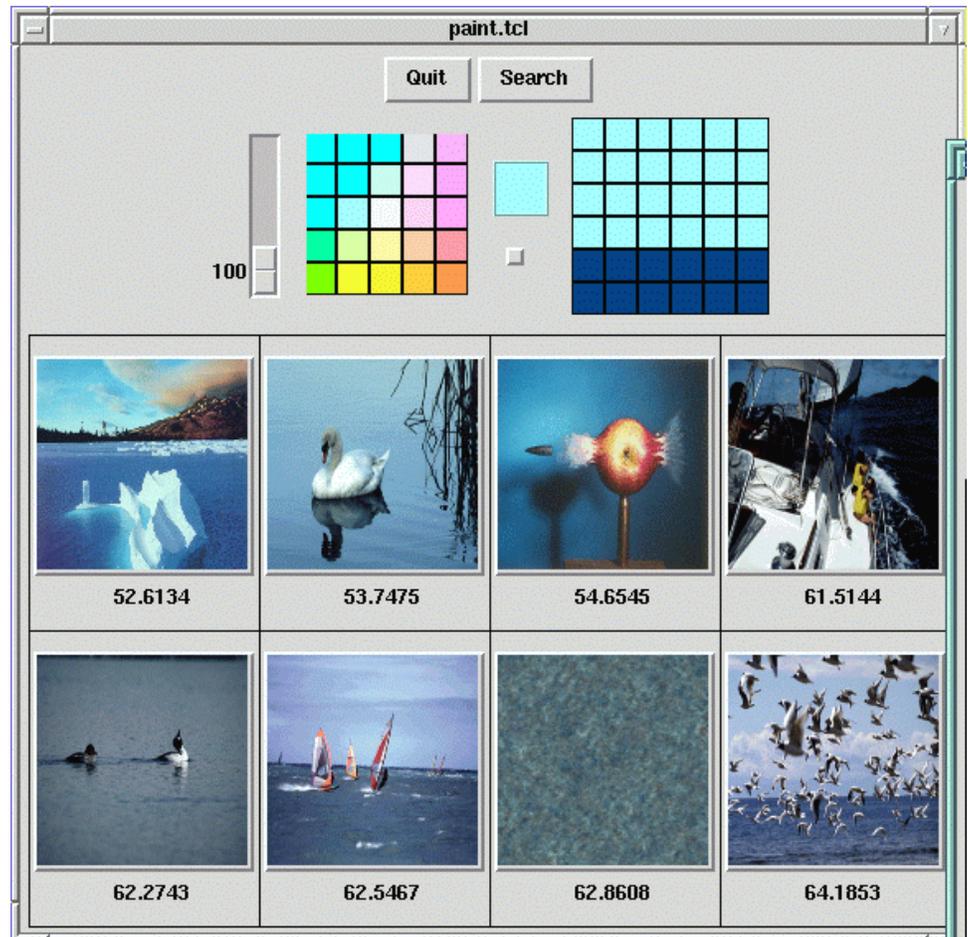


Figure 8.5: Results of an image database search in which the query is a painted grid.

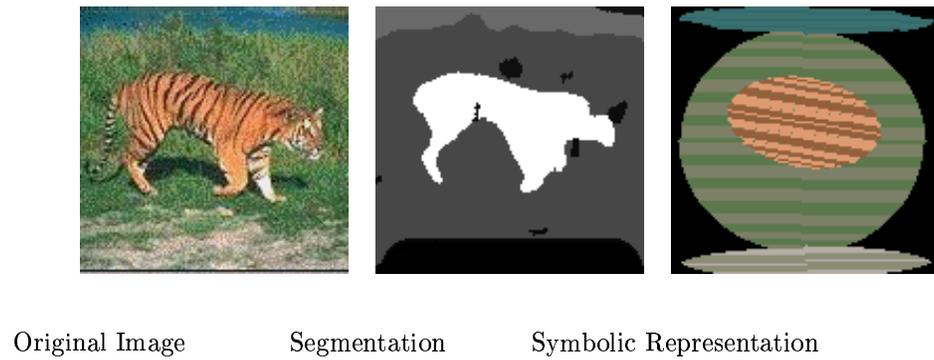


Figure 8.9: Objects and spatial relationships that can be extracted from images and used for retrieval.

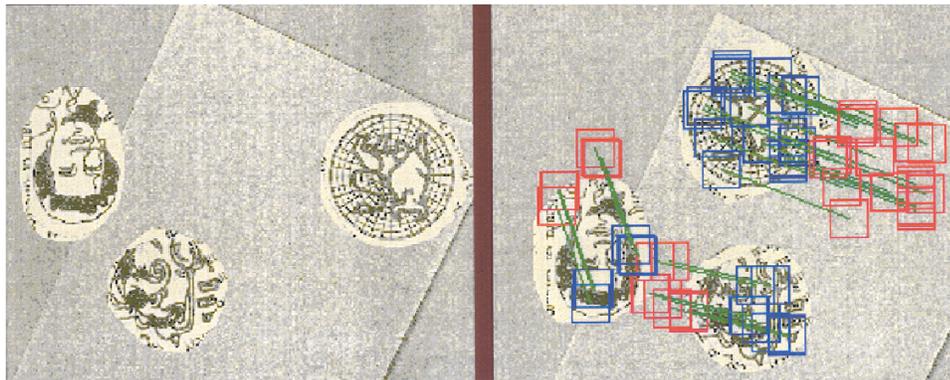


Figure 9.7: Results of applying Algorithm ??. At the left is the image at time t_1 . At the right is the image at time t_2 with the motion analysis overlaid. Red squares indicate the location of the original neighborhoods detected by the interest operator in the image at the left. Blue squares indicate the best matches to these neighborhoods in the image at the right. There are three coherent sets of motion vectors (green lines) corresponding to three moving objects. Analysis courtesy of Adam T. Clark.



Figure 10.1: Football image (left) and segmentation into regions (right). Each region is a set of connected pixels that are similar in color.



Figure 10.4: Football image (left) and $K=6$ clusters resulting from a K-means clustering procedure (right) shown as distinct gray tones. The six clusters correspond to the six main colors in the original image: dark green, medium green, dark blue, white, silver, and black.

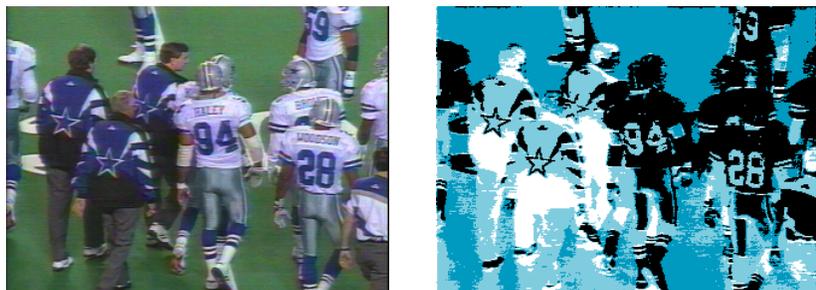


Figure 10.5: Football image (left) and $K=5$ clusters resulting from an isodata clustering procedure (right) shown as distinct gray tones. The five clusters correspond to five color groups: green, dark blue, white, silver, and black.



Figure 15.16: Rough mesh model of a dog and texture-mapped image. Courtesy of Kari Pulli.

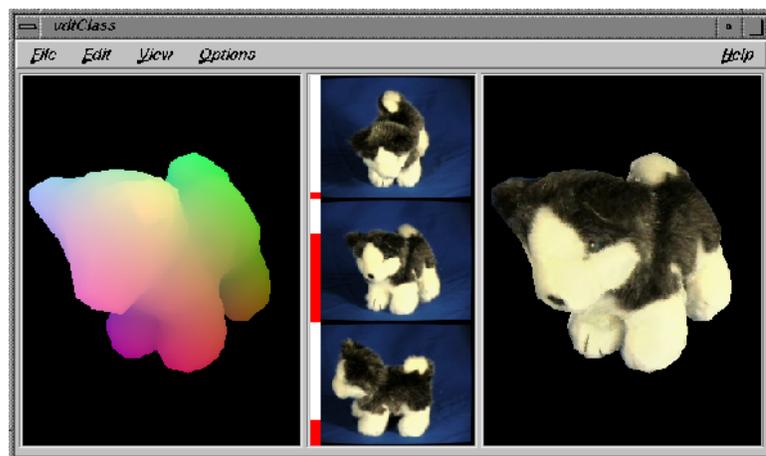


Figure 15.17: Range data of a dog model (left), three real color images from nearby viewpoints (center), and the rendered image using a weighted combination of pixels from these views (right). Courtesy of Kari Pulli.

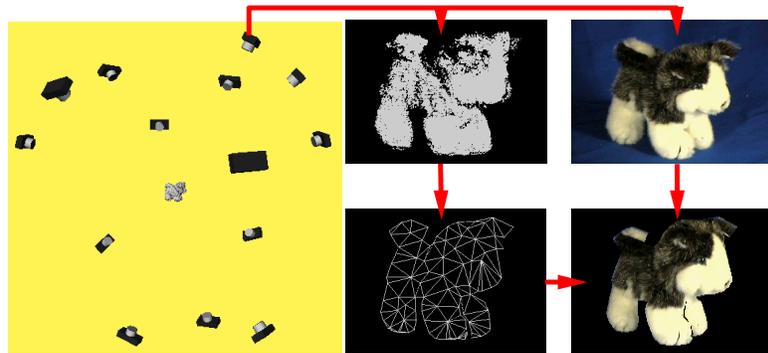


Figure 15.18: Registered range and color images from a small number of views of an object can be used to produce a high-quality rendered image, without ever constructing a full model of the object. Potential viewpoints (left), range data from one of the viewpoints (top center), color data from the same viewpoint (top right), mesh constructed from the range data (bottom center), and rendered image achieved by texture-mapping the color data onto the mesh (bottom right). Courtesy of Kari Pulli.



Figure 15.19: The same technique applied to an object for which construction of a full 3D model is nearly impossible, due to the thinness of parts of the object. Three different color images of the object (upper left), three images from a different, selected viewpoint constructed by mapping the pixels of the three original images onto the new viewpoint (lower left), and final rendered image, which is a weighted combination of the three constructed images (right). Courtesy of Kari Pulli.