Displaying Radiologic Images on Personal Computers: Image Processing and Analysis

Thurman Gillespy III and Alan H. Rowberg

This is the fourth article of our series for radiologists and imaging scientists on displaying, manipulating, and analyzing radiologic images on personal computers. Classic image processing is divided into point, area, frame, and geometric processes. Point processes change image pixel values based on the value of the pixel of interest. Histogram equalization adjusts the pixel values in the image based on the distribution of pixel values. Area processes change the pixel of interest based on the values of the surrounding pixels, known as the neighborhood. Area processes using a convolution kernel are often used as image filters. Common convolution kernels include low-frequency, high-frequency, and edge-enhancement filters. Edge enhancement can be performed with convolution kernels such as shift and difference, gradient-directional and Laplacian filters, or with nonlinear methods such as Sobel's algorithm. Frame processes mathematically combine two or more images, often for noise reduction and background subtraction. Geometric processes alter the location of pixels within the image, but usually not the pixel values. Common radiologic applications of image processing include window width and window level adjustments (point process), adaptive histogram equalization (area process), unsharp masking (area process), computed radiography image processing (combined area and point processes), digital subtraction angiography (frame and geometric processes), region of interest analysis (area process), and image rotation (geometric process). As digital imaging becomes more widespread, radiologists need to understand the image processing that is fundamental to these modalities.

Copyright © 1994 by W.B. Saunders Company

KEY WORDS: personal computer, image processing, histogram equalization, convolution, Laplacian filter, Sobel's algorithm, region of interest, unsharp masking.

IMAGE PROCESSING can be defined as the method of converting an image from one form into another. The results are usually an alteration of the pixel values or pixel locations within the image, and a visual change in the image appearance. Image processing is used for

a wide range of scientific, artistic, and commercial applications. Image processing is essential to many disciplines that use images, including diagnostic radiology.

This is the fourth and penultimate article of our series for radiologists and imaging scientists on displaying, manipulating, and analyzing radiologic images on personal computers. First, classic image processing is briefly discussed. Then the more common image-processing applications in radiology are described. A computed tomography (CT) scan of the chest (Fig 1A) is used to illustrate several of the image-processing techniques.

CLASSIC IMAGE PROCESSING

Classic image processing is conventionally categorized as point processes, area processes, frame processes, and geometric processes.^{1,2}

Point Processes

Point processes alter the value of each pixel based on the value of that pixel. Common applications include brightening and darkening an image (by adding or subtracting a fixed value to each pixel), thresholding, and histogram equalization. Point processes are usually implemented as lookup tables, either in hardware for special purpose equipment, or in software for personal computers.

Image brightening can be expressed as:

$$Q_b(i, j) = Q(i, j) + k \quad (1)$$

From the Department of Radiology, University of Washington, Seattle, WA.

Address reprint requests to Thurman Gillespy III, MD, Department of Radiology, SB-05, University of Washington, Seattle, WA 98195.

*Copyright © 1994 by W.B. Saunders Company
0897-1889/94/0702-0017\$3.00/0*

Fig 1. (A) shows CT scan of the chest and (B) shows histogram of the image. X-axis shows CT pixel values in Hounsfield units; Y-axis shows number of pixels.

where $Q(i, j)$ is the original image, $Q_b(i, j)$ is the altered image, and k is the constant value added to each pixel (Fig 2). If k is positive, the overall brightness of the image is increased, and if k is negative, image brightness decreases.

Thresholding converts all pixel values below a threshold value to black and the other pixels to white (Fig 3).

$$Q_t(i, j) = \begin{cases} 0 & \text{if } Q(i, j) < T \\ 1 & \text{if } Q(i, j) \geq T \end{cases} \quad (2)$$

where $Q_t(i, j)$ is the thresholded image, T is the threshold pixel value, 0 is a black pixel and 1 is a white pixel. Thresholding in conjunction with edge detection is often used as a preprocessing step in machine vision and other image-analysis applications.

Contrast stretching or histogram equalization is not technically a point process, but is included

Fig 2. Image brightening is shown. A constant value was added to each pixel within the rectangular region overlying the right chest.

in this section because histograms are often used to illustrate the effect of point processes. A histogram is a plot of the distribution of the pixel values of an image (Fig 1B). Histogram equalization is the process of analyzing an image to determine what densities are present, and then adjusting the display of the image so that all the gray shades are more or less equally represented (Fig 4).

The process is implemented by counting the number of pixels that are present at each image pixel value. For example, in an 8-bit image there

Fig 3. Thresholding is shown. All pixels below -720 Hounsfield units were set to black, and all other pixel values were set to white.

Fig 4. Histogram equalization is shown. (A) shows histogram equalization of CT chest (Fig 1) and (B) shows histogram after the equalization. Compare with Fig 1B, and note the contrast range of the image has been markedly reduced, and the image pixel values are more evenly distributed. X-axis shows altered CT pixel values; Y-axis shows number of pixels.

are 256 possible pixel values. Then, the computer adjusts the image pixel values so that the number of pixels for each value is approximately 1/256 the total number of pixels in the image. Although histogram equalization may increase the visibility of certain image features, the ability to make quantitative measurements of the image is lost, and the image may have a peculiar appearance.

Area Processes

Area processes alter the value of each pixel based on the values of the surrounding pixels. The surrounding pixels are known as the neighborhood.

The neighborhood is generally a two-dimensional matrix of pixels with an odd number of elements for each dimension. The pixel to be changed is at the center of this matrix. The image is processed by changing each pixel value based on the values of the surrounding neighborhood. Most area processes use the local changes in pixel values in the neighborhood to analyze the spatial frequencies in the image. Access to the spatial frequencies allows area processes to filter the image, ie, subtract or enhance certain frequencies.

The most common method of performing area processing is to use a convolution kernel. A convolution kernel is an odd-numbered rectangular matrix of numbers known as weighting coefficients (Fig 5). Each coefficient is multiplied by the pixel underlying that coefficient, and the sum of all these operations replaces the image pixel that lies at the center of the kernel. The process is repeated for each pixel in the image. Common image-processing operations performed by convolution kernels include blurring, edge enhancement, and edge detection. Some commonly used convolution kernels are shown in Fig 6.

The convolution algorithm filters images in the spatial domain. However, images can also be filtered in the frequency domain. There are a wide variety of transform functions that can

x_1	x_2	x_3	x_4	x_5	x_6	x_7			
x_8	x_9	x_{10}	x_{11}	x_{12}	x_{13}	x_{14}			
x_{15}	x_{16}	x_{17}	x_{18}	x_{19}	x_{20}	x_{21}			
x_{22}	x_{23}	x_{24}	x_{25}	x_{26}	x_{27}	x_{28}	c_1	c_2	c_3
x_{29}	x_{30}	x_{31}	x_{32}	x_{33}	x_{34}	x_{35}	c_4	c_5	c_6
x_{36}	x_{37}	x_{38}	x_{39}	x_{40}	x_{41}	x_{42}	c_7	c_8	c_9
x_{43}	x_{44}	x_{45}	x_{46}	x_{47}	x_{48}	x_{49}			
Image Matrix							Convolution Kernel		

$$\bar{x}_{25} = (c_1 \cdot x_{17}) + (c_2 \cdot x_{18}) + (c_3 \cdot x_{19}) + (c_4 \cdot x_{24}) + (c_5 \cdot x_{25}) + (c_6 \cdot x_{26}) + (c_7 \cdot x_{31}) + (c_8 \cdot x_{32}) + (c_9 \cdot x_{33})$$

$$\bar{x}_{26} = (c_1 \cdot x_{18}) + (c_2 \cdot x_{19}) \text{ etc.}$$

Fig 5. The convolution algorithm is shown. The pixel of interest (x_{25} in this example) is replaced by the output of the convolution kernel ($\bar{x}_{25}$). Each convolution coefficient ($c_1, c_2, c_3 \dots$) is multiplied by the corresponding pixel, and the sum of these products replaces the pixel of interest. The process is repeated for each pixel in the image.

Fig 6. Common convolution kernels are shown (A) shows low-pass filters. The weighting coefficients are usually expressed as fractions. However, computation is much faster if the kernel coefficients are recast as integers, and the sum of the entire matrix is divided once by the sum of weighting coefficients (the scaling factor). (B) shows blurring filters; (C), high-pass filters; (D), shift and difference filters; and (E), gradient directional filters. Kernels for the other gradient directions are listed in Baxes¹ and Lindley.² (F) shows Laplacian filters.

convert signals and images between the spatial and frequency domains. The fast Fourier transform is the most popular transform for image conversion. For a further discussion of image filtering in the frequency domain, see Gonzalez and Wintz,³ Russ,⁴ and Barnes and Lauro.⁵

Low-pass spatial filters. Low-pass filters suppress the high-frequency components of an image, but do not alter the low-frequency components. Low-pass filters are commonly used to reduce image noise. Some common low-pass convolution kernels are shown in Fig 6. The size of the kernel and the values of the weighting coefficients determine the cutoff frequency of the filter. The sum of the weighting coefficients is one. Therefore, when the filter is moved over regions where the pixel values are the same or change very little, the new pixel value is the same as the old pixel value and the low-frequency information is not affected. However, if the pixel values are changing rapidly (large high-frequency component) the kernel averages the values within the kernel, and the high-frequency components are reduced.

Blurring filters. A convolution matrix comprised of all ones is a special case of a low-pass filter sometimes referred to as an N × N blurring filter (Fig 7). To increase the blurring, the size of the matrix is increased. In some

Fig 7. 9 × 9 blurring filter is shown.

Fig 8. High-pass filter is shown.

implementations, the sum of all the pixels in the kernel is divided by the total number of pixels in the kernel, ie, by N^2 . This class of filters has important applications in radiology (see below). Another popular blurring filter is the Gaussian blur. The weighting coefficients of these kernels approximate a Gaussian distribution when viewed along any longitudinal or diagonal axis that includes the center of the kernel.

High-pass spatial filters. High-pass filters, on the other hand, enhance the high-frequency components of an image (Fig 8). The low-frequency components are unchanged, but are attenuated relative to the high-frequency components. The effect is to enhance the edges of objects, which often improves image appearance. However, image noise is also increased. Examples of some high-pass kernels are shown in Fig 6. As with low-pass kernels, the sum of the kernel coefficients is one. Therefore, when the pixel values are the same within the kernel, the image is unchanged. However, in regions of the image where pixel values are changing rapidly, the large center coefficient value and the smaller, mostly negative surrounding coefficients amplify the changes in pixel values, and the high-frequency components are increased.

Edge enhancement. This class of filters enhances the high frequencies in the image, and usually suppresses other image information. The result is a strikingly edge-enhanced image.

The resultant image is usually not visually pleasing, but may be an important first step in image recognition and feature-analysis algorithms. In contrast to high- and low-frequency filter convolution kernels, the sum of the convolution coefficients is zero. Therefore, when the pixels within the kernel have nearly the same value, the output of the kernel is zero, which markedly attenuates the low-frequency components. After edge enhancement, a thresholding operation is often performed.

Shift and difference filters detect edges by shifting the image one pixel, and then subtracting it from the original. The result is accentuation of horizontal, vertical, or both horizontal and vertical edges depending on the convolution kernel. The output of the vertical difference kernel is quite similar to the difference image produced by differential pulse code modulation technique in image compression (see equation 4 and Figure 1, D and G in the previous article on image compression in this series).⁶ Gradient directional filters use more complex convolution kernels than the shift and difference filters, and can accentuate edges in selected directions depending on the kernel (Fig 6).

Laplacian filters, on the other hand, feature an omnidirectional edge enhancement (Fig 9). Some Laplacian convolution kernels are shown in Fig 6. For a mathematical derivation of these filters, see Gonzalez and Wintz,³ Russ,⁴ and Lindley.²

Fig 9. 5×5 Laplacian filter is shown.

Edge detection: Sobel's algorithm. All edge-detection methods using convolution are linear, ie, they are derived from the sum of the first-degree products. Sobel's algorithm, on the other hand, is a nonlinear operation that uses the first-degree products to detect edges. Sobel edge enhancement is generally a better method of edge detection than the previously described techniques and is consequently widely used (Fig 10). Classically, Sobel's algorithm is implemented with two different convolutions:^{2,4}

$$X: \begin{vmatrix} -1 & 0 & 1 \\ -2 & 0 & 2 \\ -1 & 0 & 1 \end{vmatrix} \quad Y: \begin{vmatrix} 1 & 2 & 1 \\ 0 & 0 & 0 \\ -1 & -2 & -1 \end{vmatrix} \quad (3)$$

From the output of these convolutions, the magnitude and direction of the detected edges are calculated by:

$$\text{Magnitude} = \sqrt{X^2 + Y^2} \quad (4)$$

$$\text{Direction} = \arctan(Y/X) \quad (5)$$

This method is a very computationally intensive, but other implementations of the algorithm are much easier to compute.²

Median filtering. Median filtering is an area process that does not use a convolution kernel. The pixels in the neighborhood are sorted in order of pixel value. The pixel of interest is replaced by the median or mean pixel value of

the neighborhood. This filtering is useful for removing random noise from an image.

Frame Processes

Frame processes mathematically combine two or more images that are usually either spatially or temporarily related. Common applications include motion detection in security systems, quality control for assembly systems, frame averaging for noise reduction, and image background elimination. The weather map behind the weather newscaster on the local news is among the most visible examples of frame processing.

Geometric Processes

Geometric processes change the spatial arrangement of pixels within an image based on a geometric transformation. Unlike the other image-processing methods mentioned above, the intensity of the pixels is not usually altered. Applications include image registration, correction of imaging-system defects (eg, spatial aberrations) and image rotation, scaling, and translation.

RADIOLOGIC APPLICATIONS

Image processing is already an integral component of modern diagnostic imaging, and will become increasingly important in the future as the digital imaging modalities become more widely used. Some of the more popular radiologic applications of image processing are presented next using the previous discussion as a foundation.

Window Width and Window Level Control: A Point Process

The most common example of a point process in radiology is the window width and window level control for imaging modalities like CT and magnetic resonance imaging.⁵ The window/level parameters define a linear mapping of the extended range image values into 256 gray shades for printing or display at the console. The concept of the window width and window level has been previously discussed in this series.⁷

Fig 10. 5 × 5 Sobel edge-detection algorithm is shown.

Adaptive Histogram Equalization

Classic histogram equalization adjusts pixel intensity values based on the histogram of the entire image. A variant of this technique is adaptive histogram equalization (AHE), in which the intensity of each pixel is altered as a function of the intensity of the surrounding pixels (the contextual region). AHE is a locally adaptive contrast-enhancement technique that is sensitive to local changes in image contrast.^{8,9} The result is an image where objects of markedly different intensity values are simultaneously visible, eg, lung parenchyma and the mediastinum on a chest CT.

Although the histogram equalization technique does greatly increase the visibility of objects in the image, details in the dark portions of the image, such as the lung fields, often become unusually bright, and bright portions of the image, such as bone, become unusually dark. These large local changes in brightness usually are too extreme for routine radiologic imaging. Therefore, a variant of AHE, called contrast-limited adaptive histogram equalization (CLAHE) has been developed.^{8,10} In the CLAHE technique, the histogram of the contextual region is clipped to prevent any histogram bin from being larger than some percentage of the total pixels in the image. The effect is to limit the severe local changes in contrast which may be seen with AHE, and a visually more attractive image is produced.

Unsharp Masking: An Area Process

Unsharp masking is an image-processing technique originally used in radiology with analog radiographic film. In this process, a copy of the original radiograph is blurred slightly, and the original radiograph is photographically subtracted from the blurred copy in the darkroom. This combination results in a difference image that contains mostly the high-frequency details of the image, such as the edges and other small features. This difference image is then photographically added to the original image, resulting in an image with enhanced edges. This process requires meticulous work in the darkroom, with precise registration of the various images.

Unsharp masking can also be performed digitally.^{3,4,5} Analogous to the analog method, a

blurred image is subtracted from the original, creating a difference image. The difference image is then added back to the original image. The process is summarized as follows:

$$Q_{um}(i, j) = Q(i, j) + k \cdot [Q(i, j) - \bar{Q}(i, j)] \quad (6)$$

where $Q(i, j)$ is the original image, $\bar{Q}(i, j)$ is the blurred image, k is a scaling factor, and $Q_{um}(i, j)$ is the unsharp masked image. The blurred image is typically created using an $N \times N$ blurring filter (see above). Increasing the size of the $N \times N$ kernel lowers the frequency cutoff of the filter, and increasing the scaling factor increases the amount of edge enhancement that is visible in the image.

Computed Radiography (CR) Image Processing: Combined Area and Point Processes

CR images typically undergo extensive image processing before they are printed or displayed on a monitor.^{11,12} Two types of processing are performed: spatial processing (an area process) and contrast processing (a point process). The processing is performed in two steps:

Spatial processing (performed first)

$$QL(i, j) = g(Q(i, j)) \quad (7)$$

Contrast processing (performed second)

$$QV(i, j) = f(QL(i, j)) \quad (8)$$

where $Q(i, j)$ are the unprocessed pixel values from the CR image reader. The spatial processing is a variant of the unsharp mask algorithm:

$$QL(i, j) = Q(i, j) + \beta \cdot [Q(i, j) - \bar{Q}(i, j)] \quad (9)$$

$\bar{Q}(i, j)$ is the smoothed image created by the $N \times N$ blurring filter discussed above. β is a weighing factor defined as:

$$\beta = RE \cdot RT \quad (10)$$

RE is the global enhancement factor and RT (also known as the β table) is a nonlinear lookup table that reduces the amount of edge enhancement in brighter regions of the image (which reduces noise enhancement). Rank number (RN), not listed in any of the equations is the parameter that determines the size of the smoothing kernel for $\bar{Q}(i, j)$.

Contrast processing is performed using predefined lookup tables. Most of the lookup tables have a sigmoid response that mimics the charac-

Fig 11. CR contrast processing is shown. The graph shows the GT "E" curve. The curve mimics the characteristic curve of conventional radiographic film. X-axis shows input pixel values; Y-axis shows altered pixel values.

teristic curve of conventional radiographic film (Fig 11). The parameters that describe this table include the gamma table (GT), the slope of the GT curve (GA), the gamma shift (GS), and the crossover point (GC).

Adjustment of these parameters dramatically affects image appearance. At our institution, portable chest radiographs processed in the CR unit are printed with two different sets of parameters: a normal image, and an edge-enhanced image (Fig 12). We believe understanding the effect of these CR image-processing parameters on image appearance is as important as understanding the effect of film-screen combinations on the appearance of conventional radiographs.

Region of Interest Analysis: An Area Process

Region of interest (ROI) analysis is most commonly performed on CT images. The CT pixel values, usually expressed in Hounsfield units (HU), can be helpful in classifying tissue types. Common applications include determining if a lung nodule is calcified (helpful in deciding if the nodule is benign or malignant), and determining whether masses in the liver, kidneys, and other organs are cystic or solid.

ROI analysis of CT images can also be used to quantitate the density of bone and soft tissue. To obtain reasonably precise results, an external density reference must be scanned with the patient. The reference has from two to five regions of known density that are used to calibrate the HUs.

Digital Subtraction Angiography: Frame and Geometric Process

Digital subtraction angiography is probably the most common example of frame processing in radiology.⁵ In this technique, a scout or mask image is obtained of the ROI immediately before any contrast is injected (Fig 13). Then a sequence of images are obtained during intra-arterial or intravenous contrast injection. The mask image is digitally subtracted from each of the subsequent contrast images, which greatly reduces the visibility of the background structures and greatly increases the visibility of the contrast. The technique allows excellent quality angiograms to be obtained with less contrast

Fig 12. CR portable chest images are shown. (A) shows the normal processed image. Image-processing parameters: GT = E, GA = 1.0, GC = 2.0, GS = -0.2, RN = 4, RE = 0.5, RT = R. (B) shows the edge-enhanced image. Parameters: GT = A, GA = 0.9, GC = 1.5, GS = -0.2, RN = 4, RE = 5.0, RT = Q.

Fig 13. Digital subtraction angiography is shown. (A) shows the mask image. Image quality is very poor. (B) shows the contrast image. Intra-arterial contrast can be seen, but image quality remains poor. (C) shows the subtracted image. The mask image was subtracted from the contrast image. The background structures are markedly suppressed, and the contrast is much better visualized.

injected into the patient than would be possible without the technique. To improve image quality, sometimes the mask or contrast image must be shifted slightly to correct for patient motion (a geometric process).

Image Rotation and Scaling: Geometric Processes

Image rotation is among the most common examples of a geometric process in radiology. At image-review workstations that view portable chest radiographs, the chest radiographs are occasionally initially displayed “sideways” on the monitor. All such workstations use an image flip or image rotation feature to correct this problem. The more sophisticated workstations have a recall feature that will correctly

orient the image the next time it is displayed. The prevalence of this problem has led several investigators to propose automated algorithms for correcting image-orientation defects.^{13,14} Image scaling, either magnification or minification, is another geometric process commonly found in most image-display stations. Image magnification is performed to enhance the visibility of small or subtle features in the image. Image minification may be performed to allow a very large image to be displayed on a monitor or to allow multiple images to be viewed simultaneously.

ACKNOWLEDGMENT

We thank Brent Stewart, PhD, for assistance with the image processing associated with CR.

REFERENCES

1. Baxes GA: Practical Image Processing. A Practical Primer. Denver, CO, Cascade Press, 1984

2. Lindley CA: Practical Image Processing in C. New York, NY, Wiley, 1991

3. Gonzalez RC, Wintz P: Digital Imaging Processing. Reading, MA, Addison-Wesley, 1987

4. Russ JC: The Image Processing Handbook. Boca Raton, FL, CRC Press, 1992

5. Barnes GT, Lauro K: Image processing in digital radiography: Basic concepts and applications. J Digit Imaging 2:132-146, 1989

6. Gillespy T III, Rowberg AH: Displaying radiologic images on personal computers: Image storage and compression—part I. J Digit Imaging 6:197-204, 1993

7. Gillespy T III, Rowberg AH: Displaying images on personal computers. J Digit Imaging 6:151-163, 1993

8. Zimmerman JB, Cousins SB, Hartzell KM, et al: A psychophysical comparison of two methods for adaptive histogram equalization. J Digit Imaging 2:82-91, 1989

9. Rehm K, Seeley GW, Dallas WJ, et al: Design and testing of artifact-suppressed adaptive histogram equalization: A contrast-enhancement technique for display of digital chest radiographs. J Thorac Imaging 5:85-91, 1990

10. Pizer SM, Amburn EP, Austin JD, et al: Adaptive histogram equalization and its variations. Comput Vision Graphics Image Process 39:355-368, 1987

11. Sherrier RH, Chotas HG, Johnson A, et al: Image optimization in a computed-radiography/photostimulable-phosphor system. J Digit Imaging 2:212-219, 1989

12. Templeton AW, Wetzel LH, Cook LT, et al: Enhancement of storage phosphor plate images: A C-language program. J Digit Imaging 5:59-63, 1992

13. Pietka E, Huang HK: Orientation correction for chest images. J Digit Imaging 5:185-189, 1989

14. Boone JM, Seshagiri S, Steiner RM: Recognition of chest radiograph orientation for picture archiving and communications display using neural networks. J Digit Imaging 5:190-193, 1992

ERRATA

Addendum to Gillespy T III, Rowberg AH: Displaying radiologic images on personal computers: Image storage and compression—Part 2 J Digit Imaging 7:1-12, 1994

The corrected version of Table 9 follows.

Table 9. LZW Compression of the Sample Phrase "THAT THAT IS, THAT THAT IS NOT, IS NOT."

Input	Output	Input Bits*	Output Bits†	Dictionary			
				Code	Parent	Character	String‡
'T'		8					
'H'		8	9	257	'T'	'H'	'TH'
'A'	'H'	8	9	258	'H'	'A'	'HA'
'T'	'A'	8	9	259	'A'	'T'	'AT'
' '	'T'	8	9	260	'T'	' '	'T '
'TH'	' '	16	9	261	' '	'T'	'T '
'AT'	257	16	9	262	257	'A'	'THA'
' '	259	8	9	263	259	' '	'AT '
'I'	' '	8	9	264	' '	'I'	'I '
'S'	'I'	8	9	265	'I'	'A'	'IS'
' '	'S'	8	9	266	'S'	' '	'S, '
'I'	' '	16	9	267	' '	' '	' '
'S, '	264	16	9	268	264	'S'	'IS'
'T'	266	16	9	269	266	' '	'S, '
'HA'	261	16	9	270	261	'H'	'TH'
'T'	258	16	9	271	258	'T'	'HAT'
'THA'	260	24	9	272	260	'T'	'T T'
'T '	262	16	9	273	262	'T'	'THAT'
'IS'	260	16	9	274	260	'I'	'T I'
' '	265	8	9	275	265	' '	'IS '
'N'	' '	8	9	276	' '	'N'	'N '
'O'	'N'	8	9	277	'N'	'O'	'NO'
'T'	'O'	8	9	278	'O'	'T'	'OT'
' '	'T'	16	9	279	'T'	' '	'T, '
'IS '	267	24	9	280	267	'I'	'I, I'
'NO'	275	16	9	281	275	'N'	'IS N'
'T'	277	8	9	282	277	'T'	'NOT'
' '	'T'	8	9	283	'T'	' '	'T, '
< EOF >	' '		9				
	< EOF >		9				
Total		344	261				

The dictionary entries 0 to 255 are preloaded with all possible input characters. Entry 256 is the end-of-file (EOF) marker. As characters are input into the program, new entries are added to the dictionary. The coder outputs tokens, which either represent an original or new dictionary entry. In one implementation,⁷ the dictionary is a data structure with three entries: the dictionary code, the parent code, and the recently added character. The parent code can either be an original or new dictionary entry.

*Number of bits to store the uncompressed character.

†Number of bits output by the coder. A 9-bit token is used for illustration purposes.

‡The effective string represented by the dictionary code. This string is not actually stored in the dictionary per se.