

Holographic Image Representations: the Fourier Transform Method

Alfred M. Bruckstein*, Robert J. Holt, and Arun N. Netravali

Bell Laboratories, Lucent Technologies, Murray Hill, New Jersey 07974, USA

Abstract. We discuss holographic image representations. Arbitrary portions of a holographic representation enable reconstruction of the whole image, with distortions that decrease gradually with the increase of the size of the portions available. Holographic representations enable progressive refinement in image communication or retrieval tasks, with no restrictions on the order in which the data fragments (sections of the representation) are accessed or become available.

1 Introduction

From arbitrary portions of optical holograms that encode a scene, the entire scene can be reconstructed. The quality of the view recovered depends on the size of the hologram portion used, but is independent of the place from where it was cropped. A small portion of the hologram provides a blurred scene reconstruction, and as the area of the portion grows, the scene is reconstructed with more and more detail. The information on the scene is homogeneously distributed in its holographic representation, in such a way that even small portions of the hologram contain a global view, with details missing. Hence an optical hologram of a 3-D scene is a representation that makes possible a successive refinement scheme for information retrieval and transmission, insensitive to the order in which the portions of the representation become available.

In this paper we discuss two ways of representing an image holographically. These representations do not simulate the process of optical holography, yet they achieve the property of distributing the image information “uniformly” in a way that enables successive refinement insensitive to the order in which the data portions are made available to the user.

2 Holographic Fourier Representations

Suppose $I(x, y)$ is the image to be holographically represented. The idea is to regard this positive image as the *amplitude* of the Fourier Transform of its holographic representation $H(x, y)$, defined to yield a spatially uncorrelated random phase. Hence

$$H(u, v) = \mathcal{FT}^{-1}\{I(x, y)e^{jP(x, y)}\}$$

* On sabbatical from the Technion - IIT 32000, Haifa, ISRAEL.




Fig. 1. (a)–(b) Original test images of size 256×256 . (c)–(d) Original test images of size 512×512 . (a)–(d) are synthetic, with (a) and (c) having three distinct gray levels, while (b) and (d) have continuous tones.

where $P(x, y)$ is a “random phase” image so that $E[P(x, y)P(\tilde{x}, \tilde{y})] = 0$ for $(x, y) \neq (\tilde{x}, \tilde{y})$ and $P(x, y)$ is a random variable uniformly distributed over $[-\pi, \pi]$. Let us analyze why $H(u, v)$ is expected to be a holographic representation of $I(x, y)$. Here by holographic representation we mean that from an image portion $H^c(u, v)$ cropped from the complex image $H(u, v)$ we can, by 2-D Fourier Transformation, get a version of $I(x, y)$ so that the degradation is proportional to the size of $H^c(u, v)$. Clearly if $H(u, v)$ is available we get back $I(x, y)$ as the amplitude of its 2D Fourier Transform, and the same image recovery process will be applied to cropped parts of $H(u, v)$.

To see why we expect this idea to work let us look at the one-dimensional version of the above proposed holographic representation method. The discrete 1-D signal $\{I(k)\}$ for $k = 0, 1, \dots, M-1$ (where M is usually 2^n for some positive integer n) is transformed to

$$H(u) = \sum_{k=0}^{M-1} I(k)e^{j2\pi P(k)} \cdot \frac{1}{\sqrt{M}} e^{j\frac{2\pi}{M}uk} = \frac{1}{\sqrt{M}} \sum_{k=0}^{M-1} I(k)e^{j\frac{2\pi}{M}[uk+MP(k)]}$$

where $\{P(k)\}$ is a set of independent and identically distributed (i. i. d.) random numbers uniformly distributed over $[0, 1]$. We shall represent the process of cropping a portion of $H(u)$ by multiplying it with a window function $W(u)$ so

that

$$W(u) = \begin{cases} 1 & \text{for } u \in [a, a + (L - 1)] \\ 0 & \text{for } u \notin [a, a + (L - 1)] \end{cases}$$

where $a \in \{0, 1, \dots, M - L\}$ for simplicity.

Now we ask the question: what can be recovered from $H^c(u) = H(u) \cdot W(u)$ by the 1-D Fourier Transform? After some algebra, we can obtain

$$I_W(r) = \frac{1}{M} e^{-j \frac{2\pi}{M} (a + \frac{L-1}{2}) r} \sum_{k=0}^{M-1} I(k) e^{j \frac{2\pi}{M} [MP(k) + (a + \frac{L-1}{2})k]} \frac{\sin \frac{L\pi}{M} (r - k)}{\sin \frac{\pi}{M} (r - k)} .$$

This equation shows that $I_W(r)$ will be, essentially, a “low-pass filtered” version of $I(k) e^{j2\pi[P(k) + (a+L-1)k/M]}$. If the random phase factors were not present here the $I_W(r)$ would be simply a “band pass filtered” version of $I(k)$ and the holographic property would be lost, since the key low frequency components would usually be missing in this case.

The role of the random phase $P(k)$ added to $I(k)$ when generating $H(u)$ is therefore clear: since for any a , the sequences $\{\tilde{P}_a(k)\} = \{[P(k) + \frac{k}{m}(a + \frac{L-1}{2})] \bmod 1\}$ are i. i. d. uniformly over $[0, 1]$ due to the fact that the $P(k)$ are uniformly i. i. d. there, all portions of $H(u)$ of length L will yield essentially the “same” type information on the whole sequence $I(k)$.

We should consider the behavior of the coefficients

$$\phi_L(r - k) = \frac{1}{M} \frac{\sin \frac{L\pi}{M} (r - k)}{\sin \frac{\pi}{M} (r - k)}$$

for $k = 0, \pm 1, \pm 2, \dots$, to see what type of an average of $I(k)$ the value of $I_W(r)$ is. First of all we see that when $L = M$ we have

$$\phi_L(r - k) = \begin{cases} (-1)^{(M-1)[(r-k)/M]} & \text{if } r \equiv k \pmod{M} \\ 0 & \text{if } r \not\equiv k \pmod{M} \end{cases} ,$$

hence $I_W(r) = I(r) e^{j2\pi P(r)}$ as expected. As L decreases from M we realize that $I_W(r)$ will be a local average of $I(k)$, involving k 's around r with various random phase factors affecting the averaging.

Suppose $I(k) = I_0$ is constant (for all k , or a suitably large neighborhood of r). Then we have

$$I_W(r) = I_0 \sum_{k=0}^{M-1} e^{j2\pi \tilde{P}_a(k-r)} \phi_L(r - k)$$

and

$$|I_W(r)|^2 = I_0^2 \left\{ \sum_{k=0}^{M-1} \phi_L^2(r - k) + \sum_{k \neq l} \sum_{k, l=0}^{M-1} e^{j2\pi [\tilde{P}_a(k-r) - \tilde{P}_a(l-r)]} \phi_L(r - k) \phi_L(r - l) \right\} .$$

The question is: how good an estimate is $|I_W(r)|$ for the value I_0 . We rely on the following result, which is proved in [2].

Lemma 1. Let $\theta_0, \theta_1, \dots, \theta_{M-1}$ be i. i. d. random numbers uniformly chosen from $[0, 1]$, and define the random variable $V = \sum_{k=0}^{M-1} e^{j2\pi\theta_k} \varphi(k)$, where $\{\varphi(k)\}$ is a sequence of real numbers (weights). For V we have the following statistics:

$$E(V) = 0, \quad E(|V|^2) = \sum_{k=0}^{M-1} \varphi^2(k), \quad \text{and} \quad \sigma(|V|^2) = \sqrt{\sum_{k \neq l} \sum_{k,l=0}^{M-1} \varphi^2(k) \varphi^2(l)}.$$

From the lemma it follows that

$$E \left[|I_W(r)|^2 \right] = I_0^2 \sum_{k=0}^{M-1} \phi_L^2(r-k)$$

and

$$\sigma \left[|I_W(r)|^2 \right] = I_0^2 \sqrt{\sum_{k \neq l} \sum_{k,l=0}^{M-1} \phi_L^2(r-k) \phi_L^2(r-l)}.$$

If $L = M$ we get $E[|I_W(r)|^2] = I_0^2$, $\sigma[|I_W(r)|^2] = 0$, and if $L < M$, we get

$$E \left[|I_W(r)|^2 \right] = \frac{L}{M} I_0^2.$$

If $L \leq (M+1)/2$,

$$\sigma \left[|I_W(r)|^2 \right] = \frac{L}{M} \sqrt{1 - \frac{2L^2 + 1}{3LM}} I_0^2,$$

while if $(M+1)/2 \leq L \leq M$, we have

$$\sigma \left[|I_W(r)|^2 \right] = \frac{1}{M} \sqrt{\frac{(M-L)(M^2 - 5ML + 10L^2 - 1)}{3M}} I_0^2.$$

Therefore we have that for L small in comparison with M , $(M/L)^{1/2} |I_W(r)|$ is a relatively noisy estimate of I_0 with standard deviation proportional to I_0 , as in an exponential distribution. However we have the possibility to further improve our estimators by local spatial averaging, and it is exactly this factor that intervenes crucially in image representation and recovery.

So far we have assumed that $H^c(u) = H(u) \cdot W(u)$ is used to recover the holographically encoded signal. However we might have also considered recovering the signal $\{I(k)\}$ from $H(u)$ for $u \in [a, a + (L-1)]$ by computing the 1-D Fourier Transform of the sequence $\{H(a), H(a+1), \dots, H(a+(L-1))\}$ directly. In this case we can use the so-called ‘‘packing theorem’’ for Fourier Transforms [1, p. 369]. If $\mathcal{FT}\{H(a)H(a+1) \cdots H(a+(L-1))\} = I_T(r) \big|_{r=0,1,\dots,L-1}$ and $\mathcal{FT}\{H(a)H(a+1) \cdots H(a+(L-1)) \underbrace{00 \cdots 0}_{M-L}\} = I_{WS}(\nu) \big|_{\nu=0,1,\dots,M}$, then for $\nu = 0, L/M, 2L/M, \dots, (L-1)L/M$ we have

$$I_{WS}(\nu) = \sqrt{\frac{L}{M}} I_T \left(\frac{\nu L}{M} \right).$$

But we know that

$$I_{WS}(\nu) = I_W(\nu) \cdot e^{-j\frac{2\pi}{M}a\nu}$$

since $I_{WS}(\nu)$ is a (cyclic) shift of $H^c(u) = H(u) \cdot W(u)$ by a spaces to the right. Therefore we have that

$$I_T\left(\nu \frac{L}{M}\right) = \sqrt{\frac{M}{L}} I_W(\nu) e^{-j\frac{2\pi}{M}a\nu}$$

yielding

$$|I_T(\tilde{r})| = \sqrt{\frac{M}{L}} \left| I_W\left(\tilde{r} \frac{M}{L}\right) \right| \text{ for } \tilde{r} = 0, 1, 2, \dots, (L-1) .$$

Hence $|I_T(\tilde{r})|$ is a subsampling of $(M/L)^{1/2}|I_W(r)|$ at $r = 0, M/L, 2M/L, \dots, (L-1)M/L$. Since $(M/L)^{1/2}|I_W(r)|$ is an estimator of I_0 (in case of constant $\{I(k)\}$), we have that $|I_T(\tilde{r})|$ is readily an estimate of I_0 for $\tilde{r} = 0, 1, 2, \dots, (L-1)$. Therefore we can proceed either way: consider that the data provided by $\{H(u)\}$ for $u \in [a, a + (L-1)]$ is a sequence of length L and compute the magnitude of its Fourier Transform or pad $H(u)$ with the appropriate number of zeros (to get a sequence of length M) and compute the magnitude of its Fourier Transform and then multiply resulting sequence by $(M/L)^{1/2}$ to get an estimate of $I(k)$. The second approach has the apparent advantage of automatically generating the interpolation (from the sparse data) for full-sized image reconstruction, at the expense of doing larger size transform computations; however, as we shall see the interpolation provided by this method is not visually superior to pixel replication (see Fig. 2).

So far we have seen that $H(u, v)$, the complex image associated with $I(x, y)$ via the inverse Fourier Transform $\mathcal{FT}^{-1}\{I(x, y)e^{jP(x, y)}\}$, is a holographic representation of $I(x, y)$. For simplicity, we have made the arguments to prove this in the 1D case; however they extend in a straightforward manner to 2D images. This representation maps an image matrix $I(x, y)$, encoded usually by 8 bits per pixel, with a pair of matrices, the real and imaginary parts of $H(u, v)$, that must be quantized with sufficient precision to yield “good quality” glimpses of $I(x, y)$ from arbitrary portions of $H(u, v)$ and full recovery of $I(x, y)$ from the complete $H(u, v)$.

We have first tested this idea of holographic encoding of $I(x, y)$ using the computer’s double precision floating point representation for the real and imaginary parts of $H(u, v)$. For this case the precision of the recovery of $I(x, y)$ from the entire $H(u, v)$ via the Fourier Transform was down to the least significant bit of the 8-bit representation of $I(x, y)$. A 512×512 image with 3 gray levels (27–127–227) (see Fig. 1(c)) was recovered as having 6 gray levels (26, 27), (126, 127), (226, 227). When a simple uniform quantization was used to represent $H(u, v)$ with a total of 4 + 4 bits per pixel (4 bits for each the real and imaginary parts of $H(u, v)$), the variance increased to about 20 levels about each of these three gray levels. With 6 + 6 bits per pixel the variance was about 10 levels and the resulting reconstruction of $I(x, y)$ was visually almost indistinguishable

from the original. At $8 + 8$ bits per pixel the variance was about 4 levels and the reconstruction looked perfect. (Only when histogram equalization was done could one see the variability in each region having “constant” gray levels.) At each of these quantization levels we also computed the entropy of the resulting histograms for the images $\text{Re}H(u, v)$ and $\text{Im}H(u, v)$. The histograms had a “Gaussian”-like appearance and the entropies were, for example, 5 bits (for the 6-bit representations), 7 bits (for the 8-bit representations), 9 bits (for the 10-bit representations), and 15 bits (for a 20-bit representations). In view of these and other experimental findings we concluded that, with optimal quantization in the (u, v) domain, the $\text{Re}H(u, v)$ and $\text{Im}H(u, v)$ images can in fact be as reliably represented with a total of 8 bits per pixel as their straightforward uniform quantization to $8 + 8$ bits, (and this was used in all recovery experiments).

Figure 2 compares the recovery of the 256×256 test image shown in Fig. 1(b) from arbitrary portions of $H(u, v)$ of various sizes using either zero padding of $H^c(u, v)$ to an array of original size followed by taking its Fourier Transform or the direct application of the Fourier Transform followed by expansion to 256×256 via straightforward pixel replication. These results show that recovery via direct application of the Fourier Transform followed by pixel replication is quite good and the interpolations implicitly carried out via the zero padding approach clearly does not justify the added computational effort. Subsequent experiments with 512×512 images hence always used image recovery by direct Fourier Transform and pixel replication.

Figures 3(a)–(f) shows reconstruction of the synthetic images of Fig. 1(c) and (d) from arbitrary portions of $H(u, v)$ of sizes 256×256 (1/4 of original), 128×128 (1/16 of original), and 64×64 (1/64 of original). The fact that the amplitude information of a random phase image is distributed “evenly” in the Fourier Domain was previously pointed out in the context of SAR imaging, see e. g. [4]; however this idea was never proposed as an effective way to achieve distributed image representation.

3 The Holographic Sampling Method

Another idea for holographic image representation is based on sequential sampling of the image. The original 2-D array of image pixels is stacked into a sequence so that any portion of this sequence will correspond to a subsampling of the image, with samples spatially distributed as uniformly as possible. The best results were obtained with a “pseudo-random” uniformly distributed sequence, and smoothing the resulting recovered images. For details see [2, 3].

4 Conclusions

There exist many well established types of hierarchical representations of images; however it seems to us that the holographic representation as discussed herein and others that will probably be developed in the future do have some advantages in a “distributed” world.




Fig. 2. Comparison of image recovery by zero padding ((a), (c), (e)) and by using Fourier Transform directly ((b), (d), (f)) and expanding the result by pixel replication. (a)–(b) 128×128 portions. (c)–(d) 64×64 portions. (e)–(f) 32×32 portions.

We have presented here two ideas for holographic image representations. One can reconstruct a blurred version of the whole image from an arbitrary portion of these representations, with the degree of blurriness decreasing as the size of the portion increases. Holographic representations enable progressive refinement in image communication tasks, with no restrictions whatsoever on the order in which the data are received.

As we have seen from the experimental results, the pseudo-random holographic sampling representation yields better image quality, especially when small portions of the data are available. For the Fourier Domain representation however, the recovery of blurred originals is possible even when the location of the cropped portion is unavailable. Both methods are, in our opinion, viable if the aim is progressive refinement for preview purposes and the full image recovery will be carried out before actually making further use of the images. We point out that holographic representations are by no means compression schemes: their use can be in distributing the information evenly among many locations in a computer network or in situations when the image is to be broken into packets of data that may arrive to a user in a random order. For these situations the user can have a low resolution preview of the image from whatever piece of information is first available. Such holographic representation could also be used for sharing image data in a way that will make it necessary for all the




Fig. 3. Reconstruction of the 512×512 synthetic images of Fig. 3 from arbitrary portions of $H(u, v)$ of various sizes: (a), (d) size 256×256 . (b), (e) size 128×128 . (c), (f) size 64×64 .

parties involved (that own pieces of the holographic representation) to agree and collaborate in order to recover the highest quality version of the image.

5 Acknowledgment

We thank Tom Marzetta for bringing reference [4] to our attention following a lecture on holographic image representation based on [2].

References

1. R. N. Bracewell. *The Fourier Transform and its Applications, 2nd ed.* McGraw-Hill, New York, 1978.
2. A. M. Bruckstein, R. J. Holt, and A. N. Netravali. Holographic image representations. Technical Report BL0113380-961010-09TM, Bell Laboratories, Murray Hill, NJ, October 1996.
3. A. M. Bruckstein, R. J. Holt, and A. N. Netravali. Holographic image representations: the Subsampling Method. *1997 IEEE International Conference on Image Processing*, Santa Barbara, CA, October 1997.
4. D. C. Munson. and J. L. C. Sanz. Image Reconstruction from Frequency-Offset Fourier Data. *Proceedings of the IEEE*. Vol. 72, 1984.

This article was processed using the \LaTeX macro package with LLNCS style