SUPER RESOLUTION BY PHASE-GRADIENT UNRAVELLING

D. FRASER

Department of Electrical Engineering,

University of New South Wales, ADFA, Canberra, ACT, Australia

and

B. R. HUNT

Department of Electrical and Computer Engineering,

University of Arizona, Tucson, Arizona, U.S.A.

ABSTRACT. In this work, we present another proof of the theoretical existence of super resolution, under certain conditions. The proof relies on an ideal model in which a prototypical discrete image is formed by summation of many discrete pulses placed anywhere on a regular grid. If the model is then band-limited in spatial frequency, the original, grid-resolution pulse image may be reconstructed from the band-limited information. The reconstruction uses "phase-gradient unravelling" — i.e., from a very limited number of terms of the discrete Fourier transform of an image, which defines a very limited spatial-frequency band, we extract or unravel the individual phase-gradients which, together, define the original image.

1. Introduction

We offer another proof of the theoretical existence of super resolution, under certain conditions, which also provides a simple conceptual model for such a process, and a possible means for obtaining super resolution in practice — to this end, some numerical examples of the method have been tested. The proof requires an ideal model in which a prototypical discrete image is formed by summation of many discrete pulses existing anywhere on a regular grid, similar to Frieden's model [1], but not restricted to binary data. Such a model has the spatial resolution of the grid, with more or less significantly quantized pixel values if the prototype pulses have unit amplitude, but may be coincident. The question we consider is, if the model is subsequently band-limited in spatial frequency, can the original, grid-resolution pulse image be reconstructed from the band-limited information?

Consider a single, unit pulse, in a "sea" of zeros. Such an image has a very simple Fourier transform, having constant amplitude and phase gradient proportional to the distance of the pulse from the origin. A more general pulse image may be formed from the sum of such single-pulse images, and this has a Fourier transform in which the phase gradients of the single-pulse Fourier transforms are "intertwined" by the corresponding complex summation in the Fourier domain. From a limited number of terms of the discrete Fourier transform (DFT) of an image, which defines a limited spatial frequency bandwidth, we extract or unravel the individual phase-gradients which define each single-pulse image. Once unravelled, we can reconstruct an original, high grid-resolution pulse image by summation of single pulse images. As in other ill-conditioned problems, additional information is required to provide a solution, which, in this case, is the number of unit pulses which form the prototypical image.

2. Theory Based on Simultaneous Equations

The theoretical development is based on two, somewhat different approaches, of which there is only space to discuss one, here. As an example, consider a prototypical image, in one dimension,

$$f(x) = \sum_{i=1}^{n} p(x - x_i), \qquad x = 0, 1, 2, \ldots, N-1, \qquad (1)$$

$$\text{where } p(x) = \begin{bmatrix} 1 & x = 0, \\ 0 & \text{otherwise,} \end{bmatrix}$$

and where $\{x_i : i = 1, 2, \ldots, n\}$ is a set of n arbitrary, integer pulse positions, each in the range $x_i = 0, 1, 2, \ldots, N-1$.

The N-point DFT of $f(x)$, can be written

$$F(u) = \sum_{i=1}^{n} W_i^u, \qquad u = 0, 1, 2, \ldots, N-1, \qquad (2)$$

where $W_i = \exp(-j 2\pi x_i / N)$, and $j = \sqrt{1}$.

Each term W_i^u in the summation of Eq. (2) is a term in the DFT of a unit pulse, at offset x_i, which has constant amplitude and phase gradient in u of $2\pi x_i / N$, which is proportional to x_i. Eq. (2) may be viewed as a set of non-linear, simultaneous equations, to be solved for unknown W_i, given $F(u)$. E.g., when $n = 2$, and using only two terms of the DFT, we have

$$F(1) = W_1 + W_2$$
$$F(2) = W_1^2 + W_2^2. \qquad (3)$$

Solving for W_1 and W_2, given $F(1)$ and $F(2)$, we obtain

$$W_{1,2} = F(1)/2 \pm \sqrt{F(2)/2 - F(1)^2/4} \qquad (4)$$

from which the prototypical pulse positions, x_1 and x_2, are easily found. Eqs. (3) and (4) show that, given only the first two terms (after the zero frequency term) of the DFT of a one-dimensional image, we can reconstruct an original, high-resolution prototype image, having N pixels, if we have the additional information that all pixels are zero, except for two (arbitrarily positioned), unit-amplitude pixels.

We can extend this argument to cover $n = 3$, or more, prototypical pulses, and two or more dimensions, by increasing the number of equations appropriately, through the use of additional DFT coefficients, corresponding to a wider bandwidth. Unfortunately, the non-linearity of the equations rules out an analytic solution in general, so that optimizing search methods have to be used instead.

3. Results

To test our ideas, we simulated one-dimensional images, formed as in Eq. (1), with $N = 128$. For reasons of uniqueness in position, we only allow the pulses to occupy the first N/2 grid positions, but can choose exactly where to put each pulse. We carry out a DFT of the image, and then select a small band of DFT coefficients as our low-resolution "data", ignoring all the other coefficients, to begin a search. A test image,

which has the same number of unit pulses as the original, is constructed, with unit pulses placed at random grid points, initially. From the DFT of the test image, and within the chosen narrow band of coefficients, only, the magnitudes are compared with the magnitudes of the prototype image DFT coefficients, by forming a mean-square error, e^2, which is used as a cost function, to be minimized in a Monte Carlo fashion.

In Fig. 1 the result of a search for a set of unit pulses, positioned as shown on the bottom line (the final result of the search) is shown. In the top line, the random, initial position guess is shown, and successive lines show each guess whenever e^2 is reduced from a previous minimum. On the right of each line, the iteration count and mean-square error are printed, requiring, in this case, nearly 40,000 guesses to find the solution. Thus, we see the pulses gradually "marching in" towards their final resting positions. At about iteration 900, in this case, a different search strategy is adopted, in which the current, guessed pulse positions were considered likely to be fairly close to the final positions, and a local search around the current positions was applied, having a Gaussian random distribution over a small range of positions.
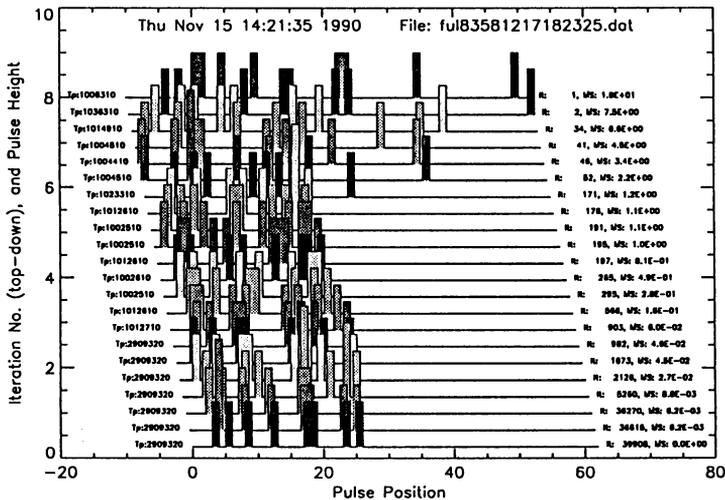
Fig. 1 Example Monte Carlo search

4. Conclusion

We have presented some preliminary results showing super resolution at work in a special case — a prototype image which can be modelled as a sum of unit pulses in a sea of zeros. Our work represents a different attempt to look at super resolution. The effect of noise has yet to be investigated, while the a priori assumption that we "know" the number of unit pulses in the prototype may or may not be useful in practice. A similar constraint in other methods is that of prior knowledge of support size, or extent of the non-zero part of a prototypical image [2].

5. References

[1] Frieden, B.R., and Zolatin, C.K.: 1986, "Monte Carlo restoration of binary objects", J. Opt. Soc. Am. A, 3, pp. 731–734.

[2] Gerchberg, R.W.: 1989, "Superresolution through error function extrapolation", IEEE Trans. on ASSP, 37, pp. 1603–1606.