IMAGE PROCESSING BY INTENSITY-DEPENDENT SPREAD (IDS)†

Tom N. Cornsweet
University of California, Irvine

SUMMARY

As retinal illuminance is lowered, the human visual system integrates the effects of photon absorptions over larger areas and longer times. A theory of the process that might underlie these changes is called Intensity-Dependent Spread (IDS). Each input point gives rise to a pattern of excitation that spreads to a region of output points, each output point delivering a signal proportional to the total excitation it sees. The unique aspect of the theory is the assumption that, although the amplitude of the excitation pattern at its center increases with input illuminance, its width decreases in such a way that its volume remains constant.

Application of this theory to image processing reveals that it displays a number of unexpected and potentially useful properties. Among them are edge enhancement and independence from scene illumination.

INTRODUCTION

During the last several years, some of my colleagues and I have been working with an interesting image processing technique called Intensity-Dependent Spread, IDS (ref. 1,2). The things I will say here are the result of my working with Jack Yellott, Steve Reuman, and Greg Reese and of discussions with a lot of other people, George Westrom, Fred Huck, and Ellie Kurrasch to name a few. I would like to discuss some of the basic properties of IDS here; some of the papers that follow this one will present specific implementations and applications of the technique.

IDS was originally developed as a theory to explain some important phenomena in human brightness perception and it has turned out to do remarkably well at relating a number of phenomena that had always before been considered quite independent of each other. But it quickly became evident that the theory had potential as a useful computer image processing algorithm too and, although I will hint at some of the relevant phenomena of human vision in this paper, the following discussion will largely be confined to IDS as an image processing technique.

†This work was partially supported by NASA contract NAS1-18468.

SOME PRELIMINARIES

The IDS model, that is, the principles underlying IDS as an image processor, can be stated very simply. However, for those who are intimate with standard image processing techniques, there are some pitfalls to understanding it that need attention. In a typical image processing procedure, for example convolving an input image with a difference of Gaussians to achieve "edge enhancement", the value at each output pixel is determined by operating on a set of corresponding input pixels, in this example applying a weighing function and then adding up the results. Although one can correctly understand the IDS procedure in that same way, I think it is much easier to understand it and to avoid pitfalls if one imagines the process as we think of optical image formation. That is, each point in the input (scene) delivers its signal (light) to some region of the output (image). This spread of signal in the image plane from a unit point in object space is called the Point Spread Function (PSF). The image can then be considered the summation of the images of all the points in object space, that is, the convolution of the PSF and an ideal image.

Now I want to introduce a new term. The PSF is the distribution of light in the image of a point of <u>unit</u> intensity. If two point sources are imaged and one source is twice as intense as the other, although it can be said that the PSF's are the same, the actual distributions of light in the two images are different, one being twice as intense as the other at every image location. To talk about IDS, we need a term that permits that differentiation. Here I will refer to the actual distribution of signal that corresponds to a particular input point as the Signal Spread, SS. The SS and the PSF for a given input point are only the same when the point happens to have unit intensity.

In an ordinary image, the volume of the SS equals the total flux emitted by the corresponding object point multiplied by a constant representing the proportion of emitted light captured by the imaging system.

I will write that in the following peculiar way:

(1)

$$V = k*Q^1$$

where V is the volume of the SS,
Q is the number of quanta emitted by the point and
K is the proportion of light captured by the imaging system.

THE MODEL

The diagram at the top of Figure 1 schematizes the IDS model. The input is represented as a distribution of values in an array of input pixels. Each input pixel delivers a signal to a network, where the signal spreads laterally. Finally, there is an array of output pixels each of which simply sums all the signals that arrive in its vicinity from the network. In our new terminology, the SSs are developed in the network and the output array sums them. Further, in the version of the IDS model to be discussed here, the SS's are everywhere positive.

If the SS's were simply, say, Gaussians whose amplitudes were proportional to the corresponding input intensities, then this model would just describe ordinary linear low-pass filtering, for example, as would result from diffraction at the pupil of the imaging system. The linear model would be appropriate if all of the energy in the output distribution were to come directly from the input, as in an optical image, or perhaps with local amplification, as is true in a photographic image or, in effect, with a standard television image. The IDS model is based on a different physical notion, that the energy at each input point <u>modulates</u> the corresponding SS, and specifically, that the input does not determine the energy in the SS but instead affects the degree to which it spreads.

Now we can state the central feature of the IDS model. Although the mathematical form of the SS is constant, for example, it is always Gaussian or conical or cylindrical, etc., its width changes inversely with input intensity. Specifically, in IDS processing, the height at the center of the SS increases when the input intensity increases and its width decreases in such a way that the volume of the SS is constant. An example for an SS of conical shape at two input intensities is shown at the bottom of Figure 1. The following equation expresses this relationship.

$$(2) \qquad V = K*Q^0$$

(Equations (1) and (2) are written this way partly to clarify an important aspect of the relationship between IDS and a linear system, but it is also meant to suggest that it would be interesting to look at the consequences of using exponents other than 1 and 0.)

That is the entire IDS model. I will just fill out two details. First, the height at the center of the SS is taken as some power function of the corresponding input strength. The simplest such function, which will be used in the following examples, has a power of one, that is, the center height is linear with input intensity. Second, although the specific details of the results are somewhat affected by the particular spread shape chosen, e.g., Gaussian <u>vs</u> cylindrical, all of the general properties I will discuss here apply for spread functions of any shape.

GENERAL CHARACTERISTICS OF IDS-PROCESSED IMAGES

Applying the IDS model or process to images produces some results that are surprising. Figures 2a and 2b summarize a group of important characteristics of IDS processing. IDS is inherently a non-linear process (because the Point Spread Function varies with local intensity, superposition is not obeyed), but the curves in Figures 2a and 2b can be interpreted as close relatives of the MTF of a linear system. Here we will call these curves Contrast Sensitivity Functions (CSFs). Consider just one curve first, say the one labeled "10" in Figure 2a. This curve shows that, at a mean intensity of 10 arbitrary units, the system acts as a bandpass filter. Therefore, for a step input the output will be a spatial transient, as plotted in Figure 3. That is, IDS does what is often called "edge enhancement". This is surprising because IDS involves no subtraction. The PSF's are everywhere positive. (If the system were linear, low frequency attenuation could only be achieved with a PSF containing some negative regions.)

Figure 2a also shows that as the mean intensity of the input changes, the CSF changes, a result that can only occur in a non-linear system. When the mean intensity increases, the entire CSF shifts toward higher spatial frequencies. Specifically, when the center peak height of the SS is linear with input intensity, the CSF shifts (on a log frequency plot) in direct proportion to the square root of the mean intensity (2).

The consequences of this shift are interesting. What the system does is automatically adjust its smoothing and spatial resolution in accordance with local photon noise. Suppose, for example, that there is a region of an optical input image that has a low mean irradiance, so that quantal fluctuations in that region render the image noisier there than in another, brighter, region. The SS's in the dark region will be larger, the CSF's there will be shifted toward lower frequencies, and each output pixel will summate signals coming from a larger input region. That is, photon detections will be summated over a larger region of the image, causing increased averaging or smoothing there. That is a good property to have, because if a region of the image is noisy, it is not possible to achieve high scene resolution there anyway. High resolution in the processing system just reveals the noise, not the details of the scene.

If, on the other hand, a region of the image has a high irradiance so that the photon statistics support high scene resolution, the IDS process automatically delivers narrow PSFs there and thus achieves high resolution.

The curves in Figure 2a plot the behavior of IDS for deterministic inputs. When the input is an optical image and the Poisson statistics of photon-matter interactions are taken into consideration, the result is as plotted in Figure 2b. At extremely low intensities the IDS process acts as a linear low-pass filter. This low-pass behavior is not exactly a consequence of the model itself, but rather will occur only when the probabilistic aspects of the input are extreme, as with photon-limited detection of extremely low light level images, and I won't discuss it further here. (See ref. 3 for a complete discussion).

Figure 4 demonstrates this property of IDS graphically. Imagine a simple scene consisting only of two adjacent regions one with a reflectance of 10% and the other of 15%, the scene being illuminated and imaged. The jagged curve at the upper left in Figure 4a is a plot of the irradiance in the image of the scene when the scene illumination is 100,000 arbitrary units, and the upper right-hand curve is the resulting IDS output. The curve on the left, the input curve, is computed assuming that the illuminating light follows Poisson statistics, as all light does, and that the sensing system noise is negligible. Thus, the jaggedness in the left curve is the result of quantal fluctuations. Some of this noise is transmitted to the output image on the right.

The pair of curves in Figure 4b show what happens when the illumination on the scene is reduced by a factor of ten. The mean image irradiances on the two sides of the edge are reduced by a factor of ten (note that the vertical axis scale is magnified by ten relative to the upper left curve) and the effect of photon noise is relatively increased (by the square root of 10). The corresponding IDS output distribution is broader but not noisier. (Note that the vertical scale of the output signals is <u>not</u> increased. The fact that the amplitude of the edge response is not changed will be discussed below.)

Moving to the curves in Figure 4c, d, e and f each successive curve shows the result of another ten-fold decrease in scene irradiance. At the lowest irradiances, individual photon detections are noticeable. Although the S/N of the input images obviously increases with decreasing scene irradiance, the noisiness of the IDS output does not. In fact, the S/N remains exactly constant for the IDS outputs, as measured either by the ratio of the mean edge response amplitude to the RMS value of the output away from the edge, or by the variance in the location of the zero crossings[1]. Thus, the IDS process yields a constant S/N for images or regions of images whose local S/N ratio varies as a consequence of quantal fluctuations.

Note that no parameters of the model were adjusted between the curves in Figure 4. With regard to the output S/N, there is only one parameter to adjust, the width of the SS at some signal input intensity. This value determines the S/N that will appear at the output.

Figure 5 shows the IDS outputs to a series of step inputs similar to those in Figure 4 but where noise is negligible. The input steps are of increasing amplitude, and any linear system will give output responses that correspondingly increase in amplitude. However, the ratios of values on the two sides of all the input steps in this figure are equal, 2:1, and the figure illustrates another important property of IDS for step inputs. The response amplitude depends exclusively on the ratio of the values across the input step.

Now imagine that the input patterns in Figure 5 are actually plots of the intensity distributions in the images of a step between two areas, one having twice the reflectance of the other, the different plots corresponding to different scene illuminations (as in Figure 4). It is then clear that, when the amplitudes of the edge responses are considered, the IDS responses to edges in a scene are independent of the level and the uniformity of the illumination on the scene. They depend only on the relationships among the reflectances in the scene. This property, independence from scene illumination, can be extremely useful when the physical properties of the surfaces in the scene are of interest. In perhaps the most important application of this property, we can show theoretically that the spectral reflectances, or more loosely the "actual colors", of objects in a scene can be determined regardless of the color of the illuminant, by applying IDS processing to each of a set of multispectral images. We are currently working on ways to exploit this IDS property in processing actual multispectral images.

A FEW SPECIFIC EXAMPLES OF IDS PROCESSING

Figure 6 illustrates the action of IDS on a television image. Because edges produce responses of equal magnitude whether in direct light or deep shadow, the output image has a much larger visual dynamic range than the unprocessed image.

An extreme case is shown in Figure 7. The input is a standard television image of a simulated space scene, using a model spacecraft and astronaut and simulating the intense shadows of space by careful baffling of the illumination. A disadvantage of IDS processing, the broadening of edge responses at low light levels, is also clearly illustrated here. Other examples of IDS processing will be given in other papers in this collection.

A MODIFICATION TO PERFORM TEMPORAL PROCESSING

In the IDS model, signals spread laterally from each input point. Suppose we add the postulate that this signal spread is not instantaneous, but rather that the signals propagate laterally with a constant velocity, as they might if they were carried by neurons, for

----

[1] These are not really zero-crossings but "base level" crossings, the base level being non-zero, dependent upon an arbitrary choice of a particular parameter of the model, and not important.

example. If the propagation velocity is taken to be very high compared with the processing rate, then all of the resulting outputs are as described above. Similarly, if the input image is stationary and the output is displayed only after the system has reached equilibrium, the results will be as above. However, if it is assumed that the lateral spreading of the signal occurs within a time scale of the same order as the time to process an image, then an interesting set of temporal properties are manifested.

Figure 8 plots temporal responses of the system when the input is a step change in the irradiance of a spatially uniform field. The different curves are for different step amplitudes. The curves suggest temporal band-pass filtering and show that, when propagation velocity is included, the temporal properties of IDS are closely analogous to its spatial properties, these curves being the temporal analog of the spatial edge responses in Figure 3. In fact, a plot of the response of the system to inputs of zero spatial frequency (spatially uniform fields) that are modulated temporally at various frequencies and with various mean irradiances looks very much like the corresponding spatial result shown in Figure 2. The system is a temporal band-pass filter that shifts toward higher frequencies as the mean irradiance increases. Thus, merely by adding the assumption that the signal spread occurs over time, the system then not only trades spatial resolution against spatial smoothing but also trades temporal resolution for temporal smoothing. That is, as light levels are reduced, the signals are automatically integrated over both larger areas and longer times.

CONCLUDING REMARKS

Certain properties of the human visual system change as the mean light level changes. In particular, as the light level is reduced, the human visual system sums the effects of detected photons over larger areas of the retina and over longer time intervals. The usefulness of that behavior in a quantum-limited detection system like the eye is clear. High system resolution in both the spatial and temporal domains is obviously useful at high light levels, but it is useless at low light levels because fine spatial and temporal detail are obscured at the input by photon statistics. To maintain a constant ability to detect an object over variations in illumination level, one must integrate over larger temporal or spatial regions, or both, as the illumination is lowered.

Intensity-Dependent Spread is an algorithm that automatically adjusts its spatial and temporal areas of integration in inverse relation to the local image irradiance in such a way that, for quantum limited detection, the S/N is constant and independent of image irradiance. The same algorithm also results in band-pass filtering and edge "enhancement", and produces responses to edges whose amplitudes are proportional to the ratios of irradiances on the two sides of the edge. It thus yields an output image of a scene that is relatively independent of the intensity and uniformity of the light illuminating the scene.

References

1. Cornsweet, T.N. "Prentice Award Lecture: A simple retinal mechanism that has complex and profound effects on perception", Amer. J. Optom. & Physiol. Optics, 62, 427 (1985).

2. Cornsweet, T. N. and Yellott, J. I. Jr., "Intensity-Dependent Spatial Summation", J. Opt. Soc. Amer. A, 2, 1769 (1985).

3. Yellott, J. I. Jr. "Photon Noise and Constant Volume Operators", J. Opt. Soc. Amer. A, 4, 2418 (1987).

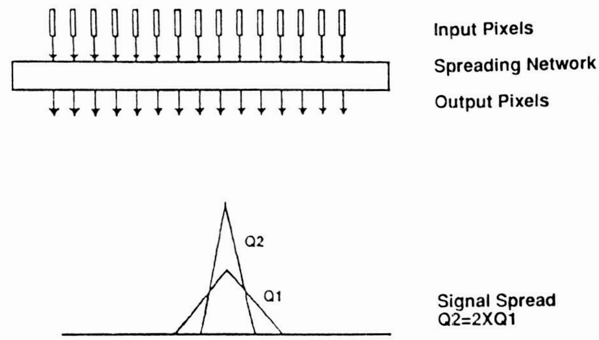

Figure 1 A schematic representation of the components of the IDS theory.

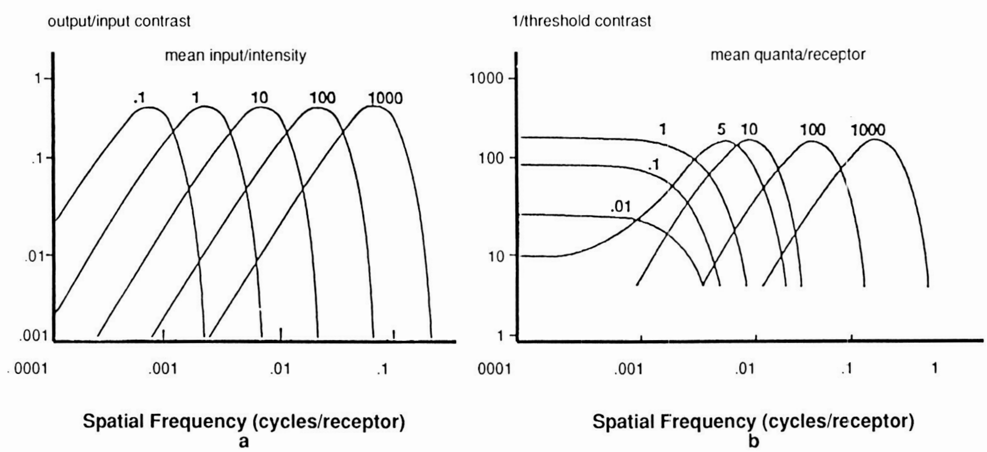

Figure 2 (a) Contrast sensitivity functions for IDS at various mean irradiances, assuming deterministic inputs. (b) Contrast sensitivity functions for IDS when photon statistics are included in the simulation. The lowest mean irradiances are such that the probability that a pixel will detect zero photons is significantly greater than zero.

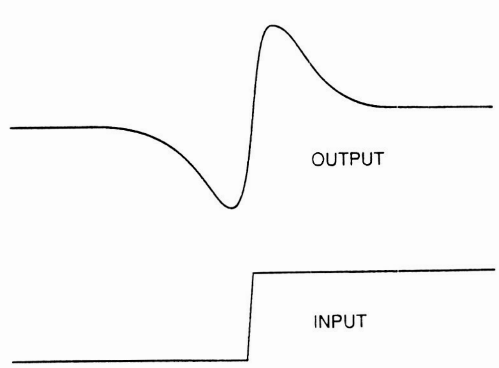

Figure 3 The IDS response to a step or edge.

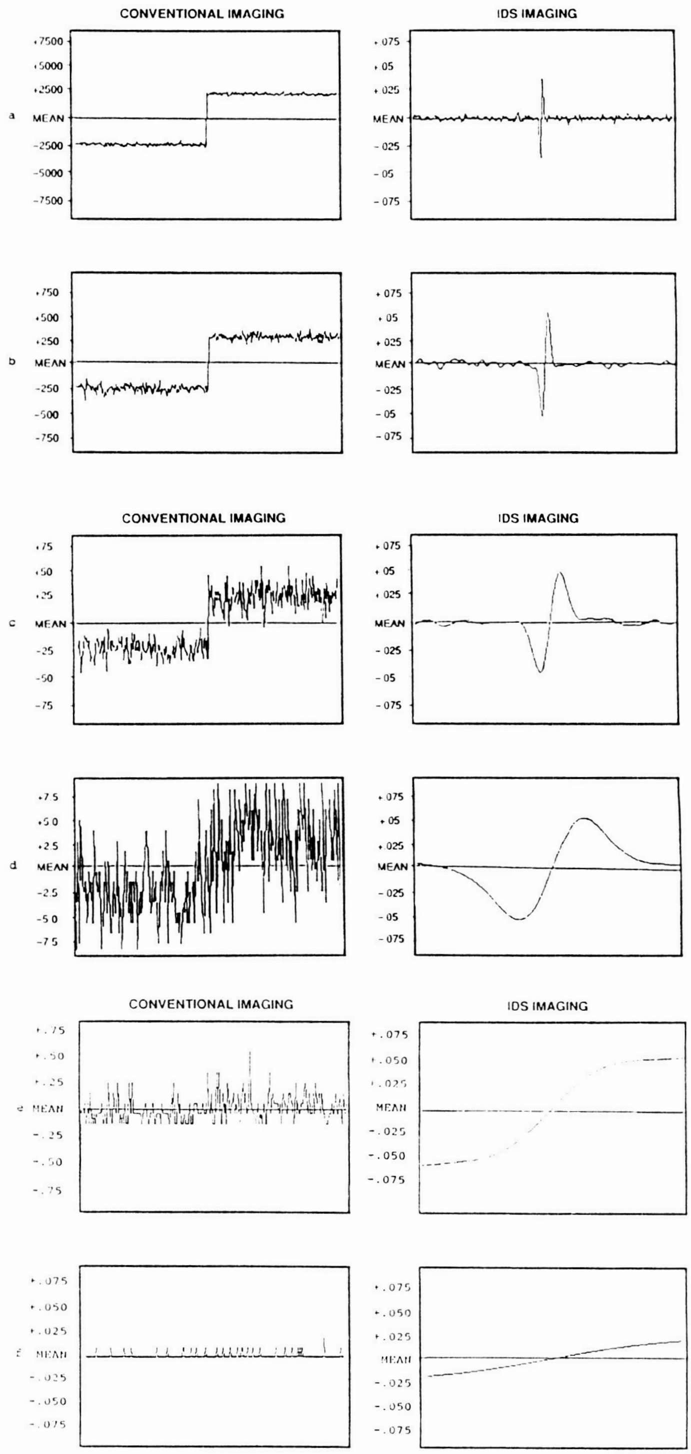

Figure 4 Each curve on the left is a plot of the relative irradiance in the image of a scene. The scene consists of two regions having reflectances of 10% and 15%. The image irradiances are computed on the basis of Poisson statistics. The curves on the right are the corresponding IDS responses. In (a), the scene irradiance is assumed to be 10,000 arbitrary units, and it is reduced by a factor of ten for each successive pair of curves.

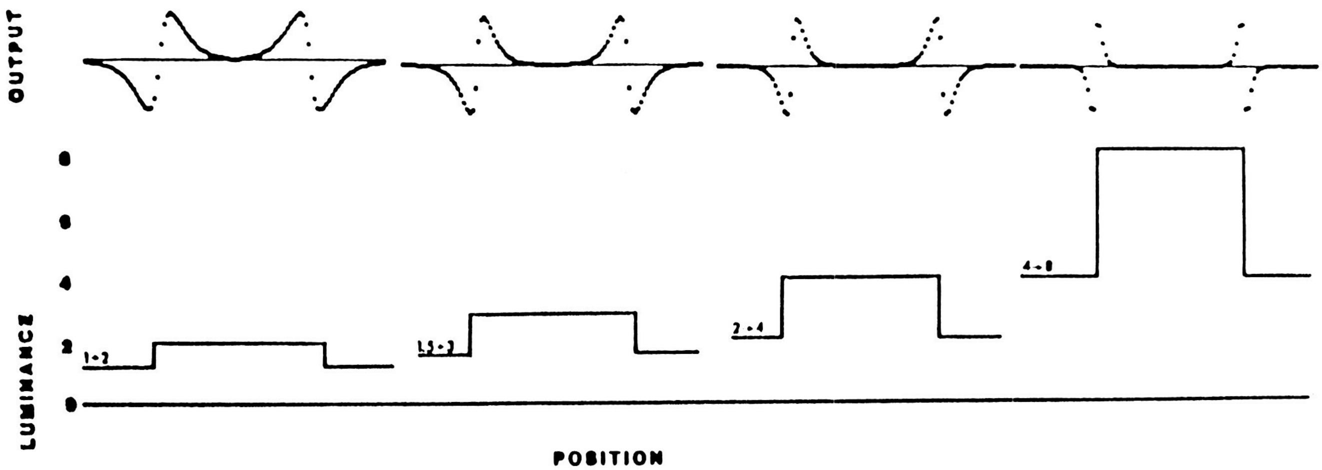

Figure 5 IDS responses to a set of deterministic step inputs, the ratio of image irradiances across the step being 2:1 in all cases.

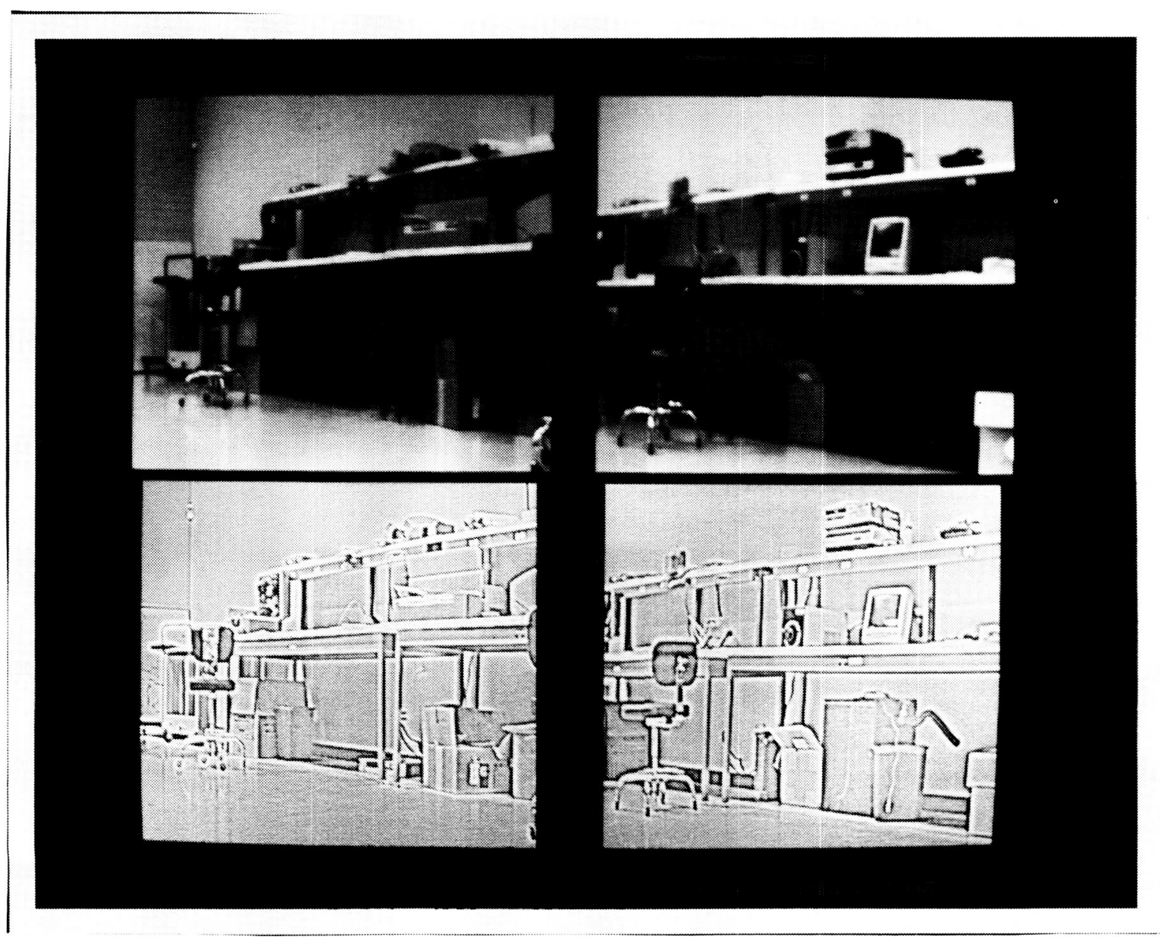

Figure 6 Two images captured by a standard television camera before and after IDS processing.

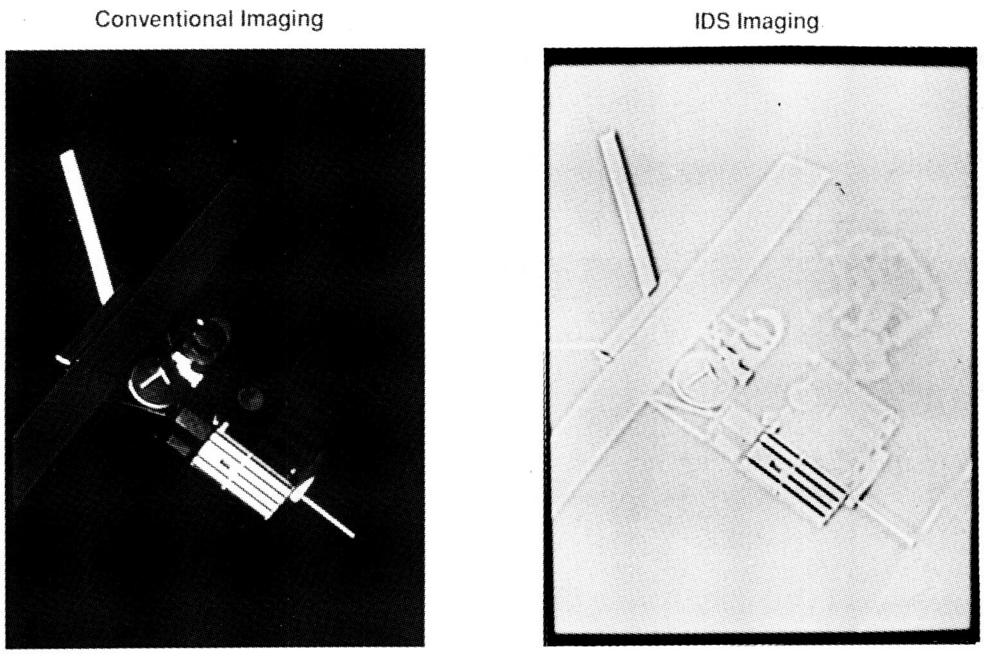

Figure 7 An image from a standard television camera of a scene simulating deep shadows in space and the result of IDS processing.

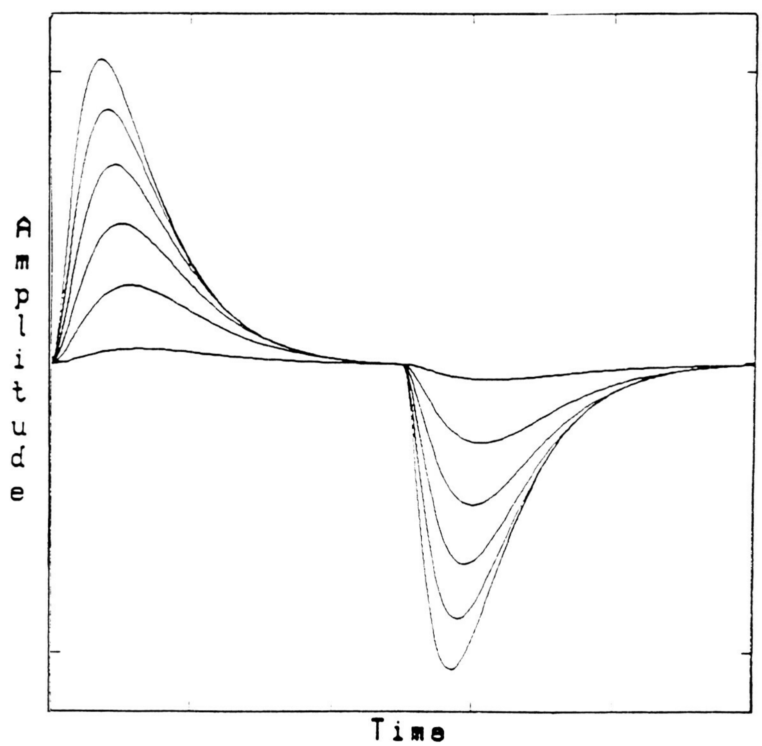

Figure 8 The responses of IDS to a spatially uniform field undergoing temporal step changes in irradiance of various amplitudes.