

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

ARTIFICIAL INTELLIGENCE LABORATORY

Artificial Intelligence

May 1975

Memo No. 333

ON VISUAL DETECTION OF LIGHT SOURCES

Shimon Ullman

ABSTRACT

The paper addresses the following problem: Given an array of light intensities obtained from some scene, find the light sources in the original scene. The following factors are discussed from the point of view of their relevance to light sources detection: The highest intensity in the scene, absolute intensity value, local and global contrast, comparison with the average intensity, and lightness computation. They are shown to be insufficient for explaining humans' ability to identify light sources in their visual field. Finally, a method for accomplishing the source detection task in the mondrian world is presented.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defence under Office of Naval Research contract N00014-70-A-0362-0003

On Visual Detection of Light Sources.

1. The Problem.

1.1 It is sometimes illustrative to think of some of the research done in Lower Level Vision as aimed at answering the following question: Suppose you are given the diagram illustrated below, which is an array of light intensities taken from some picture. The question is: what is there?

```
42 74 60 121 106 71 23 70 64 47 112 30 70 123 100 44 103 66 42 40 104
45 35 20 121 116 121 37 30 46 125 37 32 54 64 25 40 105 72 67 77 105 3
47 22 70 114 57 56 30 16 55 43 52 26 116 54 32 43 43 57 43 65 71 14 17
63 113 127 17 75 33 42 32 76 76 130 56 15 23 22 56 72 114 24 115 70 121
9 12 73 13 23 63 47 73 111 61 42 65 131 74 57 26 130 122 41 125 122 127
121 34 63 142 66 22 20 64 123 46 62 30 14 27 36 64 22 57 13 115 35 74
115 123 122 56 72 114 24 115 70 121 12 115 35 64 25 20 55 12 73 13 23 65
47 73 111 61 42 65 131 74 57 126 130 122 141 125 122 127 123 41 15 102 33
47 77 75 113 102 42 106 12 30 100 103 114 60 121 36 121 37 20 106 111 115
23 45 42 74 60 85 77 71 23 70 64 47 112 30 70 123 100 44 103 66 42 38
45 35 20 121 116 121 37 30 46 125 37 32 54 64 25 40 105 72 67 77 88 91
12 73 123 123 63 47 73 111 61 42 65 131 74 57 26 130 122 141 125 122 127
```

The correctness of the answer to the question "What is there?" is to be determined by its consistency with people's response to the same question, when presented with the corresponding picture. The main point is to emphasize that the only data available to the visual system is an array of light intensities, or perhaps a few of them for the different kinds of receptors in the eye.

In the particular problem addressed here, the general problem "What is there?" is replaced by "where are the light sources in the picture". As the answer we are after is to be consistent with human competence, it is natural to ask whether people are indeed capable of

detecting light sources in their visual fields. The answer is that in many cases they undoubtedly are. For example, in most cases it is easy to tell whether a light bulb is turned on or off. A simple demonstration of this ability is the use of brakes-light in automobiles as a warning system. This is not to say that humans can always detect light sources correctly. On the contrary: we cannot detect a weak light source if the background is too bright, e.g. stars in daytime, while on the other hand, areas which are not actual light sources are sometimes perceived as such, as in the case of the retro-reflecting shining signs on the highways. However, the interesting point, as we shall see, is not that we occasionally make erroneous judgments, but that in many situations we make the right ones.

The problem of detecting light sources has been confined here to "Achromatic Mondrians". By an Achromatic Mondrian we mean an array of rectangular shapes, of different sizes, and different levels of black, grey and white, as in picture 2. The term "Mondrians" for such arrays was used by E. Land and J. McCann, [Land 1971] due to their resemblance to the paintings by Piet Mondrian. These Mondrians serve to simplify the environment, especially by excluding colors, (Evans, in [Evans 1974] has some work on the contribution of colors to the perception of fluorescence.) and by discarding both shadings and the type of "fuzzy" light sources, shown in picture 1.

The Mondrian is composed of pieces of paper, glued together onto a white, transparent, background sheet. It is then placed on a thick piece of cardboard. By making a hole in the cardboard, and placing it above a fluorescent lamp, the area above the hole becomes a uniform light source.

A subject is then presented with the Mondrian, and asked whether he detects any

light source in it, and where. The questions we raise are, under what conditions will the subject be able to detect light sources in the Mondrian, and by what possible methods can such a task be accomplished.

2. Six possibilities, and a discussion of their insufficiency.

We shall proceed by examining the following six factors from the point of view of their relevance to light-sources detection.

1. The highest intensity in the visual field.
2. High absolute intensity value.
3. Local contrast.
4. Global contrast.
5. Intensity compared with the average intensity in the scene.
6. Lightness computation.

We shall conclude that even in the simple case of the Mondrian, these factors are not sufficient to account for the ability of human subjects to detect light sources in their visual field. The following notations are used in the sequel: The light falling on the Mondrian is called the *illumination*, and is denoted by "I". The radiant area in the Mondrian is denoted by "A"; the light transmitted through it is called the *source-intensity*, and is denoted by "L".

2.1 The highest intensity in the scene.

A perceived light source is not necessarily associated with the highest light intensity in the field. Suppose, for example, that we create an illumination gradient, and at the place where the intensity is low, we place a weak light source. The intensities can be set in such a way, that the light source will be perceivable, while the intensity graph will look like that in

figure 1:

The intensity at B is higher than the intensity at A. It is, in fact, the highest intensity in the scene. Still, only A is perceived as a light source.




Figure 1.

Thus, having the highest intensity in the scene is neither a necessary, nor is it, obviously, a sufficient condition, for being perceived as a light source.

2.2 High absolute intensity value.

A light source is not necessarily associated with an area of high absolute intensity. Against a dark enough background a firefly is perceived as a light source, although the absolute light intensity coming from it is very low. We shall resume this subject of absolute intensities after describing the results of two simple experiments which will help to elucidate our point, as well as serve as an introduction to the next section.

Experiment 1: If, at a given illumination I , we begin with a very low source-intensity and gradually increase it, the following happens. At first, the source is not detectable. When it grows stronger, there is a range of uncertainty, and then, finally, it becomes prominent. The range of uncertainty is bounded by two "thresholds": the lower threshold is the source intensity for which the subject suspects for the first time there might be a light source there; the upper threshold is the source intensity at which he becomes sure. This distinction is important when carrying out experiments with a "just noticeable" light source, as it specifies

the experimental condition more precisely.

Experiment 2: If the source intensity is set just above the upper threshold, and we gradually increase the illumination, the reverse process takes place: via the upper threshold, the range of uncertainty, and the lower threshold, to a situation where the source is no longer detectable.

In this last experiment where L was constant and I increased, the total intensity coming from area A increased, while at the same time it ceased to appear as a light source. The conclusion is that it was not the high intensity value at A that made it appear radiant. Rather, it seems that the major role was played by the ratio between the intensity at A and that of the surroundings. We shall turn therefore to consider this ratio.

2.3 Contrast:

A variable that seems suitable to explain the above results, is the contrast. The contrast between intensity I_1 and intensity I_2 can be defined as $(I_1 - I_2) / (I_1 + I_2)$. As this contrast is a monotonic function of I_1/I_2 , (assuming I_1/I_2 is positive) high contrast simply corresponds to I_1 being "many times" greater than I_2 .

We can now explain the experiments described above in terms of contrast. Suppose we have an illumination $I = 100$ units on some arbitrary scale. $L = 100$ units, and instead of a Mondrian we have a uniform surface with reflectance = 0.5 (that is, it reflects 50% of the incident light.) The light intensity coming from A will be 100 (source intensity) + 100×0.5 (reflection), that is, 150 units. From the surroundings, we get only the 50 units of reflected light, so that the ratio is 3:1, or the contrast is 0.5. If the illumination is raised to 1000 units, the new intensities ratio between A and the surroundings will be only 6:5.

Although the light intensity at A is higher this time, the contrast is much lower, only 1/11.

A plausible conclusion from the discussion so far is that whenever the contrast between a stimulus A and its surroundings exceeds a certain value, the former is perceived as a light source. There is a kind of "rationale" to such a conclusion: Surfaces in nature do not usually have reflectivity values approaching the extremes of 100 or 0. A very high contrast value is thus not likely to be achieved by reflectance changes alone; rather, it might indicate the presence of a light source.

A question of interest at this point is "just how high is high". Namely, how high are the contrasts found in natural scenes without light sources, and what contrast is needed for a light source to be just noticeable (that is, just above the lower threshold). These values were measured in a scene containing pieces of a very white paper, a very dark paper, a strip of shadow, and a light source at the lower threshold intensity. The measurements were performed using the A.I. Vidicon, from which the real intensity values are difficult to obtain. The following results should serve, therefore, more as indications of the relative magnitudes than as accurate absolute values.

The following ratios of intensity values were computed:

White to black ratio: 6-9 White to shadow ratio: 13-17

Source to black ratio: 11-15 Source to shadow ratio: 19-23

In contrast with these high values, a light source was also perceivable in a picture where the ratio of intensities was in no place higher than 1.3. Thus, the perception of a light source cannot be identified with that of a high enough contrast. While a high enough contrast seems, as far as I have been able to determine, sufficient to induce the perception of

a light source, it is not a necessary condition. In other words, the presence of a light source can be "deduced" from factor(s) other than contrast. The rest of the paper is an attempt to find such a factor. What makes the question intriguing is the fact that there does not seem to be an obvious candidate for the task.

We proceed by discussing some of the more immediate candidates for a solution, and show their inadequacy. They are: global as opposed to local contrast, average illumination, and lightness computation.

2.4 Global versus local contrast:

Consider the situation in picture 2. The source intensity was set at the lower threshold. Picture 3 is of the same setting, only this time with a dark area surrounding the light source. The contrast between the source and its immediate surroundings is thus multiplied by 6-9. This change in contrast did not, however, make the light source more noticeable. The following claim can be raised here: the change in contrast between pictures 2 and 3 is only in the immediate contrast, that is, between the light source and its immediate surroundings. The global contrast, namely between the light source and the darkest area in the field, had however not been effected. I therefore repeated this experiment, this time changing the global surroundings as well, and with similar results. Using the above distinction it can now be stated, that light sources are sometimes detectable when both the global and the local contrasts were low, therefore neither is a necessary condition for the source detection.

2.5 Average illumination:

Another factor that had been considered is the influence of changes in the average

illumination. Will the thresholds of detection become lower, if we use, for example, a darker Mondrian (while maintaining both global and immediate contrasts fixed)? No such influence has been detected.

2.6 Lightness computation:

Finally, let me turn to a brief analysis of the above results in terms of lightness to see whether this approach provides us with the key to the detection of low-contrast sources. The computation of lightness from intensities (See [Land 1971], [Horn 1974]) involves the separation of sharp intensity changes from gradual ones. Thus, an intensity graph like the one shown in figure 2A, will be decomposed into a sharp-changes component (figure 2B) and a gradual-change component (figure 2C).



Figure 2A

Figure 2B

Figure 2C

One might try to use this decomposition to explain the fact that A is perceived as a light source, although the intensity at B is higher: if the lightness contrast is computed instead of the intensity contrast, then area A obviously gets the highest value. Examination of the experimental data reveals, however, that the computation of contrast via lightness, does not help solving the basic problem: contrast is still not a necessary condition. The

visual system somehow distinguishes between possible interpretations of the scene: a contrast created by light-and-dark surfaces on the one hand, and a contrast created by the presence of a light source on the other. To make the problem of "more than one possible interpretation" clearer, consider the following intensity array:

```

44 43 43 42 42 41 41 40 40 110 109 108 107 106 105 104 103 102 101 100 50 49 48 47 46
44 43 43 42 42 41 41 40 40 110 109 108 107 106 105 104 103 102 101 100 50 49 48 47 46
44 43 43 42 42 41 41 40 40 110 109 108 107 106 105 104 103 102 101 100 50 49 48 47 46
44 43 43 42 42 41 41 40 40 110 109 108 107 106 105 104 103 102 101 100 50 49 48 47 46

```

The corresponding lightness matrix discards slow changes and therefore will look like this:

```

30 30 30 30 30 30 30 30 30 100 100 100 100 100 100 100 100 100 100 50 50 50 50 50
30 30 30 30 30 30 30 30 30 100 100 100 100 100 100 100 100 100 100 50 50 50 50 50
30 30 30 30 30 30 30 30 30 100 100 100 100 100 100 100 100 100 100 50 50 50 50 50
30 30 30 30 30 30 30 30 30 100 100 100 100 100 100 100 100 100 100 50 50 50 50 50

```

In both cases, there is more than a single possible interpretation. One interpretation is that the areas simply have different reflectivities. Thus, the central area can be normalized to 1.0, and the left and right areas assigned the values of 0.3 and 0.5 respectively. A different interpretation is that the right area has a reflectivity of 1.0, the left - 0.6, and the center is a light source.

To be sure, if we knew the real reflectance of one area, we could have determined (at least in the case of Mondrians) all the reflectance values in the picture, and then we could have also assigned the label "light source" to areas of reflectance greater than one. The only trouble with this method is that there does not seem to exist a way of determining real reflectance values. Still, in many cases the visual system is able to make the right interpretation of the scene. In a picture rather similar to the array presented above, but with even less contrast, a light source had been perceived. That is, the visual system

managed to somehow pick (correctly) the light-source interpretation. How can this be done?

The following section proposes a method.

3. The proposed method.

3.1 General description of the method.

Roughly speaking, the proposed method is the following:

Given two adjacent areas, compute both their intensity-ratio and their gradient-ratio, and compare the two. If the ratios are not equal, one of the areas is a light source.

It is based on the following two observations:

1. All that is needed for the detection of a light source, is the correct values of the reflectance ratios in the scene.
2. In many cases this real reflectance ratio can be computed, even in the presence of a light source, by comparing *intensity gradients*.

Let me discuss each in turn.

1. Using the value of (r_0/r_1) :

The preceding section seemed to imply that a correct normalization of the reflectances is needed in order to detect light sources, by identifying them with areas of reflectance greater than 1. It turns out, however, that we don't need that much.

Consider two adjacent areas: area no. 0 and area no. 1, and suppose that we somehow know the reflectance ratio r_0/r_1 . Assuming the illumination is a continuous function of the position, we can then determine whether one of the above areas is a light source. The continuity of I implies that the illumination on the two sides of the borderline between the

areas is approximately equal, if the measurements are taken close enough to the borderline. From this equality it follows that if neither of the areas is a light source, the intensity ratio should equal the reflectance ratio. Hence, if the ratio values do not agree, we can deduce the presence of a light source and, as it turns out, we can also compute the actual source intensity.

2. Computing (r_0/r_1) from gradients:

As had been mentioned above, the only values permitted in the computation are the intensity values at different points in the visual field. The way used both by Land and Horn to determine reflectance ratios from these intensities is basically the following: Find the points of sharp changes in the intensity distribution, like point A in figure 3.

Then conclude that the ratio r_0/r_1 is equal to the ratio e_0/e_1 when e_0 is measured to the left of the "jump" A,

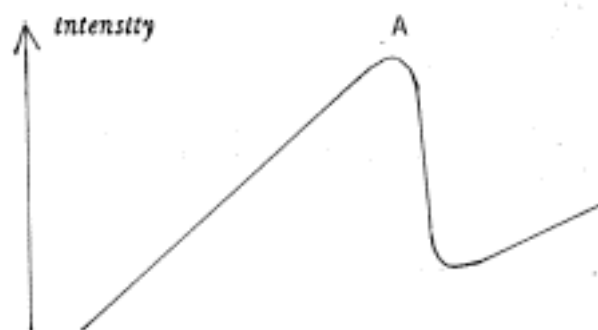


Figure 3

and e_1 is measured to the right of it. This computation, however, presupposes the existence of no light sources in the picture. For otherwise it might be the case that, for example, $r_0 = r_1$, but there is a light source at area 0 which creates the intensity jump at A. That is to say, in order to compute the reflectances ratio we need first to discover the light sources, and in order to discover the light sources we need the reflectance ratios! Here is a way out of this circle. Suppose, (and this is usually the case), that the illumination is not absolutely uniform,

but has a gradient which over some area is more or less linear, say a gradient of 100 intensity units to an inch. When such an illumination is reflected from a surface of reflectance K , not only is the overall intensity reduced by a factor of K : the gradient will be K times the original one as well.

In figure 4, the reflectances are taken to be $1/2$ in area 0 and $1/4$ in area 1.

Note that not only is $I_A/I_B = 2$, but S_A/S_B also is equal to 2; where S_A is the slope of the intensity graph measured from A to the left, and S_B is the slope measured from B to the right.




Figure 4

If we now add a uniform light source to area 0, the ratio I_A/I_B will change, but the ratio S_A/S_B will remain the same, and still equal to r_0/r_1 . S_A and S_B can, of course, be computed from the intensity values alone.

Thus, we can use the gradient ratio to compute the reflectance ratio, and then use this ratio as shown in the preceding section to test for the presence of a light source. Next, we shall describe how the computation is actually implemented.

3.2 The "Source Operator" and its implementation.

Few notations: as before, I denotes the illumination, L - the source intensity (the source is assumed to be in area 0), S is the gradient, r denotes the reflectance and e stands

for the intensity of the light reaching the eye. Note that e is the only parameter we are allowed to measure. The discussion so far shows that from the obvious relation:

$$(1) \quad e_0 = (I_0 \circ r_0) + L$$

we can deduce that under "smooth" illumination:

$$(2) \quad L = e_0 - e_1 \circ (S_0/S_1)$$

if the points 0 and 1, where the measurements are taken, are close enough to each other.

However, "real life" measurements tend to look more like figure 5.

In this case we can no longer suppose that I_A is equal to I_B , as they are too far apart. Rather, it is:

$$I_B = I_A + Kd$$




Figure 5

Where K is the intensity gradient in intensity units per unit distance, and d is the distance of B from A . Starting with the underlying equation (1), $L = e_0 - r_0 \circ I_0$, we now get:

$$\begin{aligned} L &= e_0 - r_0 \circ (I_1 - Kd) = e_0 - I_1 r_0 - r_0 Kd \\ &= e_0 - e_1 \circ (r_0/r_1) + S_0 d \quad (\text{Since } r_0 K = S_0) \end{aligned}$$

This last equation, with the slopes replacing the reflectances:

$$(S) \quad L = e_0 - e_1 \circ (S_0/S_1) + S_0 d$$

will be referred to as the "S-Operator"

It is a rather straightforward task to implement the S-Operator and run it on actual intensity arrays. Several details have to be added, mainly to cope with the problems of noise and insufficient resolution of the Vidicon. For example, the right-slope at point i cannot be computed simply by, say, $I(i+1) - I(i)$; rather some averaging is needed. In the program this average is taken over a distance which varies with the situation. Another example is the treatment of irregular points, whose intensities are exceptionally higher or lower than their surroundings. It so happens that in most cases they are lower, due perhaps to dust on the lenses.

Figure 6 at the end of this section shows the intensity distribution obtained from a scene which contained a light source to the right of point A. Figure 7 shows the corresponding output of the S-Operator. As negative S-values are meaningless in the particular algorithm that had been used, the output can be taken as the maximum of 0 and the computed S-values, as in figure 8. The "Source value" is the highest at the real source borderline. Its relative magnitude is enough to allow the setting of some threshold, above which an area can be labeled as a light source. In the general case, this threshold has to depend on the overall intensity. An S-value of 100 units, which is not significant when the intensity is around 5000 units and with a noise level of 500, is significant indeed when the intensity is around 100 and the noise level around 10 units. This fact contributes to the phenomenon that a light source, which is detectable at low illuminations, becomes undetectable at higher ones.

A final remark: The above method raises the possibility of detecting a "dark light source". An area with an average intensity of, say, 0.8 that of its surroundings, might be

"declared" as a light source, if the light gradient there is only, say, 0.4 the gradient of the surroundings. Is it possible to construct a situation in which an area with light intensities lower than the surroundings will be perceived as a light source? Although some experiments have been carried out, I don't yet know the answer to this interesting question.

References

- Evans, R. M. "The Perception of Color" John Wiley & Sons,
N.Y. London Sydney Toronto 1974.
- Horn, Berthold K. P. "Determining Lightness from an Image"
Computer Graphics and Image Processing Vo. 3, 4 1974.
- Land, E. H. and McCann, J. J. "Lightness Theory". Journal
of the Optical Society, Vol. 61, 1971.

FIGURE 6




FIGURE 7

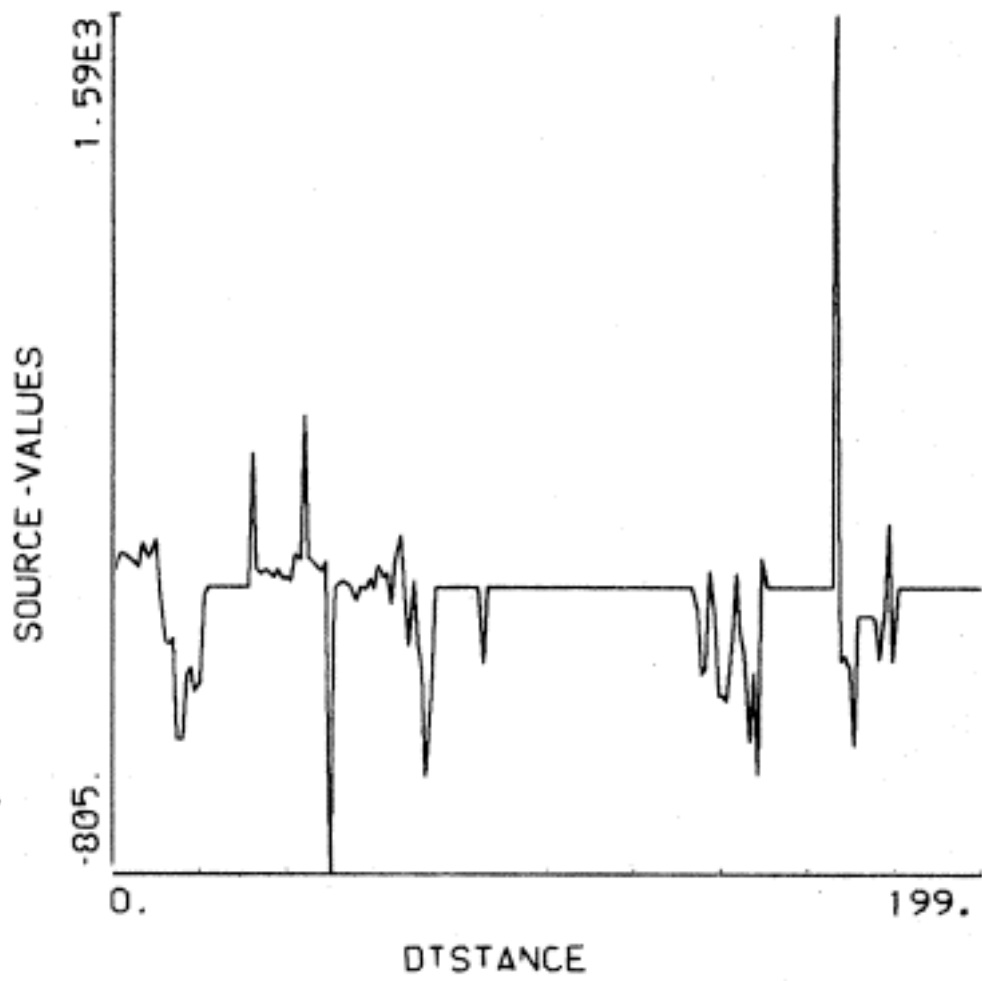





FIGURE 8





1



2



3