

Convexity-based Visual Camouflage Breaking

Ariel Tankus and Yehezkel Yeshurun

School of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel.

E-mail: {arielt,hezy}@math.tau.ac.il

Supported by grants from: Minerva Minkowski center for geometry, Israel Academy of Science for Geometric Computing, and the Moscona fund.

Camouflage is frequently used by animals and humans (usually for military purposes) in order to conceal objects from visual surveillance or inspection. Most camouflage methods are based on superpositioning multiple edges on the object that is supposed to be hidden, such that its familiar contours and texture are masked. In this work, we present an operator, (D_{arg}), that is applied directly to the intensity image in order to detect 3D smooth convex (or equivalently: concave) objects. The operator maximally responds to a local intensity configuration that corresponds to curved 3D objects, and thus, is used to detect curved objects on a relatively flat background, regardless of image edges, contours and texture. In that regard, we show that a typical camouflage found in some animal species, seems to be a "counter measure" taken against detection that might be based on our method. Detection by D_{arg} is shown to be very robust, from both theoretic considerations and practical examples of real-life images. As a part of the camouflage breaking demonstration, D_{arg} , which is non-edge-based, is compared with a representative edge-based operator. Better performance is maintained by D_{arg} for both animal and military camouflage breaking.

Key Words: convexity detection, regions of interest, camouflage breaking, counter shading.

1. INTRODUCTION

“Camouflage is an attempt to obscure the signature of a target and also to match its background” [1].

The goal of this paper is to detect 3D convex or concave objects under strong camouflage. We suggest an operator (“ D_{arg} ”), which is applied directly to the intensity function. D_{arg} maximally responds to the typical intensity signatures of smooth 3D convex or concave patches in the objects. The operator is not limited to any particular light source or reflectance function, as a large range of examples demonstrate. D_{arg} implicitly takes advantage of the 3D structure of objects, but does not attempt to restore the three dimensional scene. Although prior knowledge or experience may aid one in breaking camouflage (e.g. [2]), our operator is context-free; its only a priori assumption about the target is its being three dimensional and convex (or concave). The purpose of D_{arg} is detection of convex or concave subjects under harsh camouflage conditions or in highly cluttered scenes.

D_{arg} is robust to changes in *illumination*, *scale*, and *orientation*. The invariance theorem underlying the theoretic explanation for this phenomenon (proof supplied) is accompanied by real-life images. The theorem states that D_{arg} is invariant under any derivable and strongly monotonically increasing transformation of the intensity function. The robustness and invariance characterizing D_{arg} make it suitable for camouflage breaking, even for camouflages that might mislead a human viewer (an example of this kind of camouflage is Fig. 1(a)).


The application of the operator to real-life images demands a relatively short running time, and its robustness leads to reliable results. The reliability of D_{arg} in detection of camouflaged targets is displayed in a comparison between D_{arg} and a representative edge-based detection method: radial symmetry (see [3]). Radial symmetry is a representative edge-based method, as it was proved to generalize several other methods of interest point detection. The difficulty of camouflage breaking by edge-based methods is demonstrated in Fig. 1, where edge distribution of a military camouflage resembles the texture of the surroundings. The large number of strong edges in the scene and the camouflage clothes distract edge-based operators from the main subject.

After a short survey of related work (section 2), we define the suggested convexity-based operator: D_{arg} (section 3). Subsection 3.2 is of particular importance for understanding the behavior of D_{arg} , as it describes the goal of the operator and the way to accomplish this goal in an intuitive manner. Section 4 is a rigorous mathematical characterization of the behavior of D_{arg} . Section 5 demonstrates the robustness of D_{arg} to illumination, scale and orientation changes in real-life images; the robustness is based on the theoretic claims of section 4. Section 6 introduces biological evidence for camouflage breaking using convexity detection, based on Thayer’s principle of counter-shading. Section 7 compares the performance of the suggested convexity-based operator with that of a representative edge-based method in camouflage breaking. Several types of camouflage are examined in the comparison. Conclusions appear in section 8.

2. RELATED WORK



a. Original image.

b. Edge map ($\|\nabla I\|$).c. Radial symmetry ($r=30$).

d. Detection by the peaks of Radial Symmetry.

FIG. 1. Failure of edge-based methods. A highly textured background and camouflage clothes (a) leads to a total failure of edge-based methods. The edge map (b) is densely covered with edges, so an algorithm receiving the edge map as its input (c) fails to discriminate the 3D object from the background (d).

The work related to visual camouflage could be roughly divided into two: camouflage assessment and design, and camouflage breaking.

In the camouflage assessment literature, one may find [1], which assesses the distinctness of camouflaged patterns in a set of images. The set of images is paired, and the observer is asked to choose which of the two images possesses a target pattern that is more distinct from the background. The results of these comparative observations show high correlation with a second order metric of pattern distinctness, which is based on a model of image texture. [4] also assess the strength of a camouflaged target signature based on human perception. They combine the results of human quantification of differences in camouflage effectiveness in a semi-automatic (i.e, a human-in-the-loop) camouflage assessment system. [5] describes a system for designing camouflage patterns based on natural terrain reflectance data. [6] designs camouflage patterns based on Fourier techniques and biologically based techniques. [7] describes aircraft camouflage design research during World War II.

Despite the ongoing research, only little has been said in the computer vision literature on visual camouflage breaking: [8] considers the issue of camouflage breaking for moving targets by stereopsis, and reaches the conclusion that stereo does not greatly enhance motion detection in camouflage. [9] measures motion by interpreting object edges as differential mapping singularities. This allows motion detection even in the presence of sudden changes of target direction and use of camouflage. [10] proposes a computational model of visual moving image filtering. [11] suggests a 3D model-based aircraft recognition scheme, based on low level matching of image segments and segments of the projection of the 3D model, extracted from one or more views. The recognition system copes also with aircraft camouflage. [12] take the spectral feature as the basis for target recognition, attempting to distinguish color camouflage from green vegetation background. Breaking camouflage using light polarization has been suggested in [13]; biological evidence for polarization-based camouflage breaking in the cuttlefish was introduced in [14].

3. Y_{arg} : AN OPERATOR FOR DETECTION OF CONVEX PATCHES

We next define the operator we propose for detection of three dimensional objects having smooth convex and/or concave patches.

3.1. Defining the Argument of Gradient

Let us estimate the *gradient map* of image $I(x, y)$ by:

$$\nabla I(x, y) \approx ([D_\sigma(x) G_\sigma(y)] * I(x, y), [G_\sigma(x) D_\sigma(y)] * I(x, y))$$

where $G_\sigma(t)$ is the one dimensional Gaussian with zero mean and standard deviation σ , and $D_\sigma(t)$ is the derivative of this Gaussian. We turn the Cartesian representation of the intensity gradient into a *polar* representation. In contrast with edge-based methods, which are usually concerned with the gradient *magnitude*, we are more interested in the gradient *argument*. The argument (usually marked by $\theta(x, y)$), is defined by:

$$\theta(x, y) = \arg(\nabla I(x, y)) = \arctan \left(\frac{\partial}{\partial y} I(x, y) , \frac{\partial}{\partial x} I(x, y) \right)$$

where the two dimensional arc tangent is defined by:

$$\arctan(y, x) = \begin{cases} \arctan(\frac{y}{x}), & \text{if } x \geq 0 \\ \arctan(\frac{y}{x}) + \pi, & \text{if } x < 0, y \geq 0 \\ \arctan(\frac{y}{x}) - \pi, & \text{if } x < 0, y < 0 \end{cases}$$

and the one dimensional $\arctan(t)$ is the inverse function of $\tan(t)$ so that: $\arctan(t) : [-\infty, \infty] \mapsto [-\frac{\pi}{2}, \frac{\pi}{2}]$. The proposed convexity detector, which we mark by: Y_{arg} , is the y -derivative of the argument map:

$$Y_{arg} = \frac{\partial}{\partial y} \theta(x, y) \approx [G_{\sigma}(x) D_{\sigma}(y)] * \theta(x, y)$$

3.2. Intuition for Y_{arg}

In the next chapter we provide a mathematical description of the operator and prove its attributes. It might be useful, however, to first examine the intuition behind it.

- **What Does Y_{arg} Detect?**

Y_{arg} detects the zero-crossings of the gradient argument. This stems from the last step of the gradient argument calculation: the two-dimensional arc-tangent function. The arc-tangent function is discontinuous at the negative part of the x -axis; therefore its y -derivative approaches infinity there. In other words, Y_{arg} approaches infinity at the negative part of the x -axis of the arctan, when this axis is being crossed. This limit reveals the zero-crossings of the gradient argument.

- **Why Detect Zero-Crossings of the Gradient Argument?**

Y_{arg} detects zero-crossings of the gradient argument of the intensity function $I(x, y)$. The existence of zero-crossings of the gradient argument enforces a certain range of values on the gradient argument (trivially, values near zero). Considering the intensity function $I(x, y)$ as a surface in \mathbb{R}^3 , the gradient argument is the azimuth of a normal to the surface. Therefore, a range of values of the gradient argument means a certain range of azimuths of the normal to the intensity surface. This enforces a certain structure on the intensity surface itself.

The next section (Sect. 3.3) would characterize the structure of the intensity surface as either a paraboloidal structure or any derivable strongly monotonically increasing transformation of a paraboloidal structure (Fig. 2). Since paraboloids are arbitrarily curved surfaces, they can be used as a local approximation of 3D convex or concave surfaces (Recall, that our input is discrete, and the continuous functions are only an approximation!). The detected intensity surface patches are therefore those exhibiting 3D convex or concave structure. The convexity is three dimensional, because this is the convexity of the intensity surface $I(x, y)$ (= 2D surface in \mathbb{R}^3 ; Fig. 3(b)), and *not* convexity of contours (= 1D surface in \mathbb{R}^2 ; Fig. 3(a)). This 3D convexity of the intensity surface is characteristic of intensity surfaces emanating from smooth 3D convex bodies. For example: the image of a smooth 3D convex Lambertian surface under a point light source is a convex intensity function.




FIG. 2. (a) Paraboloidal intensity function: $I(x, y) = 100x^2 + 300y^2$. (b) Gradient argument of (a). Discontinuity ray at the negative x -axis. (c) Y_{arg} of (a) ($= \frac{\partial}{\partial y}$ of (b)). (d) Rotation of (a) (90° c.c.w.), calculation of gradient argument, and inverse rotation. (e) Rotation of (a) (90° c.c.w.), calculation of Y_{arg} , and inverse rotation. (f) Response of D_{arg} , the isotropic operator.




FIG. 3. 3D vs. 2D convexity. (a) 2D convexity: A contour is a 1D surface in \mathbb{R}^2 . (b) 3D convexity: A paraboloid is a 2D surface in \mathbb{R}^3 .

- **Summary**

We detect the zero-crossings of the gradient argument by detecting the infinite response of Y_{arg} at the negative x -axis (of the arctan). These zero-crossings occur where the intensity surface is 3D convex or concave. Convex smooth 3D objects (e.g, Lambertian surfaces under point light source) usually produce 3D convex intensity surfaces. Thus, detection of the infinite responses of Y_{arg} results in detection of domains of the intensity surface which characterize 3D smooth convex or concave subjects.

3.3. Y_{arg} Response to Paraboloids

The projection of concave and convex objects can be estimated by paraboloids (Fig. 2(a)), since paraboloids are arbitrarily curved surfaces (see [15]). Our mathematical formulation refers to a general paraboloid of the form:

$$f(x, y) = a(x - \epsilon)^2 + b(y - \eta)^2$$

where $a > 0$, $b > 0$ are constants, and (ϵ, η) is the center of the paraboloid (Fig. 2(a)). The first order derivatives of the paraboloid are: $\frac{\partial}{\partial x}f(x, y) = 2a(x - \epsilon)$ and $\frac{\partial}{\partial y}f(x, y) = 2b(y - \eta)$. The gradient argument is therefore:

$$\theta(x, y) = \arctan(b(y - \eta), a(x - \epsilon))$$

(see Fig. 2(b)). Deriving it with respect to y yields:

$$\frac{\partial}{\partial y}\theta(x, y) = \frac{ab(x - \epsilon)}{a^2(x - \epsilon)^2 + b^2(y - \eta)^2} \quad (1)$$

However, this derivative exists in the whole plane **except for the ray**:

$$\{(x, y) \mid y = \eta \text{ and } x \leq \epsilon\} \quad (2)$$

Pay attention, that the derivative of the 2D arctan function along this ray is infinite, even though the derivative at the rest of the plane (as expressed in (1)) is continuous. At a first glance, Computer Vision applications tend to “smooth out” such an infinite derivative by employing the algebraic expression of (1) (which is continuous at the whole plane except for the origin) instead of the correct mathematical derivative. In contrast with this approach, we would follow the rigorous definition of derivative, and see that the infinite derivative of the discontinuity ray could be highly advantageous. In fact, this discontinuity ray and its infinite derivative are the basis for the suggested approach.

At the ray (2), $\theta(x, y)$ has a first order discontinuity (in the y -direction), so its derivative there approaches infinity. The fact that for a paraboloid, $\frac{\partial}{\partial y}\theta(x, y) \rightarrow \infty$ at the negative ray of the x -axis, while continuous at the rest of the plane can be clearly seen in Fig. 2(c) (we define our coordinate system at the horizontal and vertical axes of the paraboloid).

The reader should note, that the argument of the gradient is being used in Computer Vision for a long time. Hough transform [16], for example, uses the argument

of the gradient to reduce the space of parameters when searching straight lines or circles in the image. Image improvement by argument-based methods is also known, and can be found in [17]. Recognition of local symmetries in graylevel images employs the argument of a convolution kernel which is complex valued [18]. Nonetheless, the novel idea in the suggested operator is not the argument of the intensity gradient, but rather the usage of the *discontinuity ray* formed by the argument of the intensity gradient. The key idea for the detection scheme we describe lies in this discontinuity ray, which is detected by derivation of the gradient argument, to receive a strong response to that ray. Looking for zero-crossings in the argument of the intensity gradient is very stable.

3.4. D_{arg} : The Isotropic Variant

The strong response of Y_{arg} at the negative part of the x -axis appears when $\theta(x, y)$ changes sharply from high values (approx. π) to low values (approx. $-\pi$). This need not happen only on paraboloidal intensity functions: cropping a strip around the negative part of the x -axis from the parabolic intensity function would still produce high values, since $\frac{\partial}{\partial y}\theta(x, y) \rightarrow \infty$ there. The intensity function of the cropped image exhibit strong convexity in the y -direction, and therefore leads to the strong Y_{arg} reaction. Similarly, a horizontal *cylinder* would cause a strong Y_{arg} reaction. A vertical cylinder, on the other hand, would cause a very weak reaction, as its horizontal axis is flat. Figures 4(a),(b) demonstrate this dependence of Y_{arg} on the convexity orientation: Y_{arg} does not detect the vertical cylinder. The reaction to the oblique cylinder is weaker than to the horizontal cylinder.

In order to avoid Y_{arg} dependence on convexity orientation, we define an isotropic operator; i.e., an operator that would strongly response to *all* convexity orientations. A general way of doing so would be to rotate the original image by $\pi - \alpha$ degrees counterclockwise, calculate Y_{arg} for the rotated image, and rotate the result back to the original angle (by $\alpha - \pi$ degrees counterclockwise). We refer to this process as calculating the α_{arg} of the image. Figures 2(d),(e) describe such a rotation of the argument of the gradient and the Y_{arg} of the paraboloidal intensity function, respectively. Fig. 2(e) can be referred to as the α_{arg} for $\alpha = 90^\circ$. The α_{arg} of the cylinders image appears in Fig. 4(b)-(e). We approximate the isotropic operator: $\int_{\alpha=0^\circ}^{360^\circ} (\alpha_{arg})d\alpha$ by: $D_{arg} = \sum_{\alpha=0^\circ, 90^\circ, 180^\circ, 270^\circ} \alpha_{arg}$. Figure 2(f) shows the D_{arg} of the paraboloidal intensity function. As expected, a strong response appears in all axes. The D_{arg} of the cylinders image strongly responds to all cylinders in all orientations (Fig. 4(f)).

4. PROPERTIES OF D_{arg}

We next consider several properties of Y_{arg} . By definition, the same properties characterize D_{arg} too.

Planar objects of constant albedo form linear intensity functions, and are usually of little interest (e.g., walls). It can be easily shown, that:

- Y_{arg} has zero response to planar objects.




FIG. 4. (a) Four cylinders (ray-tracing of a three dimensional model). (b) Y_{arg}^2 of (a). Can be also referred to as α_{arg}^2 for $\alpha = 180^\circ$. (c)-(e) The α_{arg}^2 of the cylinders image in (a). Below each image appears the appropriate α . The dependence on the cylinder orientation is obvious. (f) D_{arg}^2 of (a). D_{arg} strongly reacts to all cylinders, independent of their orientation.

- Theoretically, having a continuous image, edges of planar objects produce a finite response, while paraboloids, an infinite response. Practically, the response to edges of planar objects is smaller than the response to paraboloids.
- Linear dependence on scale.

$$D_{arg}^{I(sx,sy)}(x,y) = s \cdot D_{arg}^{I(x,y)}(sx,sy)$$

The D_{arg} of the scaled image ($I(sx, sy)$) at point (x, y) is the scaled D_{arg} of the original image multiplied by the scale factor: s .

The linear dependence on scale holds when one deals either with continuous images (in theory), or when the subsampling is dense enough to faithfully represent the original image following scaling. The choice of a particular standard deviation (σ) for the calculation of derivatives affects the working scale in similar manner. However, when one examines image locations where $D_{arg} \rightarrow \infty$, a change of σ will only affect the *size* of the domain where the theoretically-infinite response should occur, but not the relative strength of the response. Theoretically, this response would be infinite, regardless of the value of σ . In practice, this response would be higher than those of theoretically-finite domains, no matter what the value of σ is. Thus, the infinite response reduces the effect of change in standard deviation σ or scale, and allows the usage of a constant value of σ for a large range of object sizes (as later demonstrations will show).

It follows from these properties, that the response of Y_{arg} to 2D-convexity that is, convexity of contours in the plane, is relatively weak. The stronger response of Y_{arg} is to 3D-convexity (convexity of surfaces in the 3D space), for example, paraboloids.

We now present an important invariance property of Y_{arg} , followed by a practical discussion and a demonstration from real-life scenes.

THEOREM 4.1.

Let $f(x, y)$ [the original intensity function] be a derivable function at each pixel (x_0, y_0) with respect to x and y .

Let $T(z)$ [the transform] be a function derivable at point $z_0 = f(x_0, y_0)$, whose derivative there is positive in the strong sense. Define the composite function by: $g(x, y) = T(f(x, y))$ [the transformed intensity function].

In this case, the y -derivatives of the gradient arguments of $f(x, y)$ and $g(x, y)$ at point (x_0, y_0) are identical:

$$\frac{\partial \theta_g(x_0, y_0)}{\partial y} = \frac{\partial \theta_f(x_0, y_0)}{\partial y}$$

Proof. By the *chain rule*, the composite function: $g(x, y) = T(f(x, y))$ is derivable with respect to both x and y at point (x_0, y_0) , and its derivatives are:

$$\begin{aligned} g_x(x_0, y_0) &= T'(f(x_0, y_0))f_x(x_0, y_0) \\ g_y(x_0, y_0) &= T'(f(x_0, y_0))f_y(x_0, y_0) \end{aligned}$$

Let $f^0 = f(x_0, y_0)$, $f_x^0 = f_x(x_0, y_0)$, $f_y^0 = f_y(x_0, y_0)$. The argument of the gradient at point (x_0, y_0) can be written as:

$$\theta_g(x_0, y_0) = \arctan(T'(f^0)f_y^0, T'(f^0)f_x^0)$$

Since we have required that $T'(f^0) > 0$, the point $(T'(f^0)f_x^0, T'(f^0)f_y^0)$ lies in the same quarter of the plane as point (f_x^0, f_y^0) . It follows that:

$$\theta_g(x_0, y_0) = \arctan(T'(f^0)f_y^0, T'(f^0)f_x^0) = \arctan(f_y^0, f_x^0) = \theta_f(x_0, y_0)$$

The last equation states that the argument of the intensity gradient is invariant under the transformation T . Deriving the gradient argument with respect to y preserves this invariance:

$$\frac{\partial \theta_g(x_0, y_0)}{\partial y} = \frac{\partial \theta_f(x_0, y_0)}{\partial y}$$

■

Let us rephrase theorem 4.1 in the following manner:

Y_{arg} is invariant under any derivable strongly monotonically increasing transformation of the intensity function.

The practical meaning of the theorem is that Y_{arg} is invariant, for example, under linear transformations, positive powers (where $f(x, y) > 0$), logarithm, and exponent. Y_{arg} is also invariant under compositions and linear combinations (with positive coefficients) of these functions, since such combinations are also derivable and strongly monotonically increasing. The functions mentioned above and their combinations are common in image processing for lighting improvement. This implies that Y_{arg} is robust to a large variety of lighting conditions. Figure 5 demonstrates Y_{arg} invariance to $\log(\log(I))$ and $\exp(\exp(I))$ in a real-life scene.

Taking Y_{arg} invariants into account, the suggested model is not only a paraboloidal intensity function detector, but also a detector of any derivable strongly monotonically increasing transformation of paraboloids. This stability makes Y_{arg} particularly attractive for usage in various scenes where the environment is unknown before hand.

5. D_{arg} ROBUSTNESS DEMONSTRATION

D_{arg} is robust to the following three factors: illumination, scale, and orientation. The robustness to the last two factors is gained mainly due to the fact that $\frac{\partial}{\partial y}\theta(x, y) \rightarrow \infty$ at the negative x -axis of paraboloids, which is a very stable feature. Scale and orientation variations preserve this approach to infinity. Robustness to illumination changes has been proved in detail in Sect. 4. Further demonstration of the robustness of the operator can be found in [19].

Robustness to Illumination. We have proved in Sect. 4, that D_{arg} is invariant to any transformation of the intensity function which is derivable and strongly monotonically increasing. As discussed there, the practical meaning of this invariance




FIG. 5. Invariance to derivable strongly monotonically increasing transformation of the intensity function. **Top row:** The original image $I(x, y)$ is compared to $\log(\log(I(x, y)))$ and $\exp(\exp(I(x, y)))$. D_{arg} is invariant under \log and \exp . **Bottom row:** D_{arg} . The similarity between the D_{arg} of the original image and the D_{arg} 's of the transformed images is obvious.

is robustness to illumination variations. Let us examine different directions of illuminations in a real-life scene. In Fig. 6, a single point light source illuminates the scene. The light source is placed in seven positions which are equidistant from the subject, to form horizontal angles of -90° , -60° , -30° , 0° , 30° , 60° , 90° , with respect to the line connecting the subject and the camera. In each of the images, the regions of the highest D_{arg} values are those of the inverted funnel.

Robustness to Scale. Figure 7 (left) demonstrates D_{arg} robustness to scale. It should be noted that throughout this article the same dimensions for the convolution windows were used: 2×2 pixels window for the first-order derivation, and 14×14 pixels window for the derivation of $\theta(x, y)$. No a-priori knowledge of the scale of the detected objects is used. In Fig. 7 (left), a vase is photographed in 6 different scales. Each image is accompanied by its D_{arg}^2 . The vase is detected by D_{arg} in each of the scales. The smallest appearance of the vase takes approximately: 24×30 pixels, while the largest, approximately: 60×76 pixels. This means that the image of the vase is zoomed approximately by factor 2.5 (in each axis).

Robustness to Orientation. The D_{arg} operator is isotropic by definition: it has been generated to answer the problem of Y_{arg} preference of objects with maximal convexity in certain orientations (see Sect. 3.4). Figure 7 (right) shows the robustness of D_{arg} to the orientation of the maximal convexity of the object. In this figure, we have taken a threshold of 70% of the maximal squared reaction of D_{arg} . The biggest blob (in terms of number of pixels) in the thresholded image is used as our detection. In all orientations the flashlight has been correctly and consistently detected.

Threshold Sensitivity Analysis. In order to show that a detection scheme is robust, it is not enough to show that the basic operator itself is robust, but one

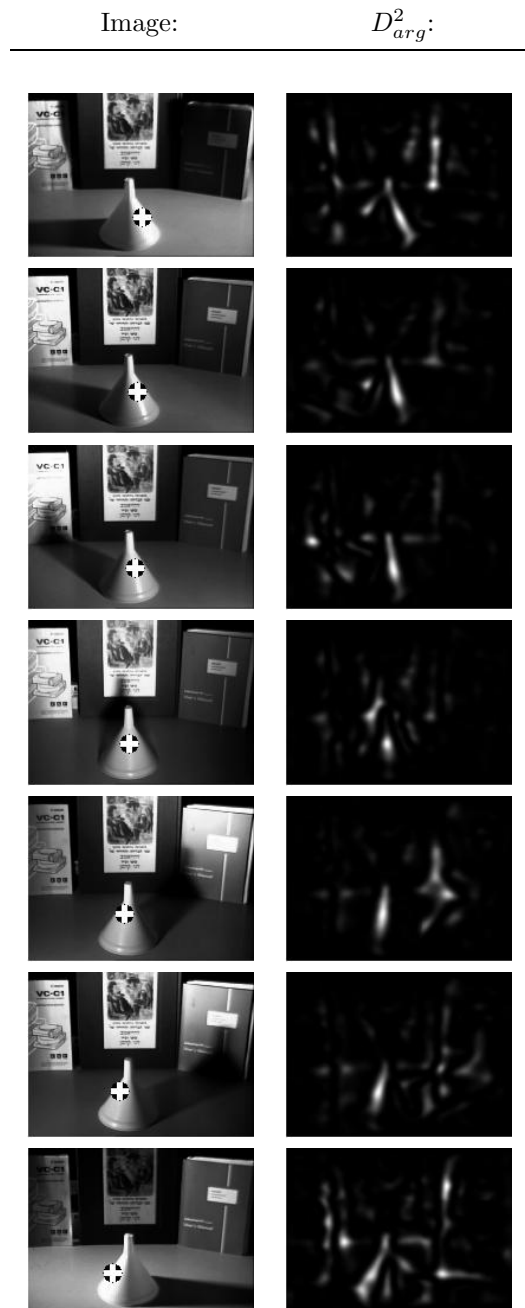


FIG. 6. Robustness to illumination. D_{arg} strongly reacts to the funnel, and is robust to changes in lighting direction. Each row corresponds to an illumination direction with an azimuth of -90° , -60° , -30° , 0° , 30° , 60° , or 90° , respectively. The detection is performed by 70% thresholding on D_{arg}^2 .




FIG. 7. **Left:** Robustness to scale changes—a vase in 6 different scales. The largest appearance of the vase is 2.5 times larger than its smallest appearance (in each axis). **Right:** Robustness to orientation variations. The flashlight changes its orientation from vertical to horizontal. D_{arg} strongly reacts to the cylindric flashlight. The detection of the flashlight is independent of orientation. The detected region is the biggest blob in the 70% thresholded D_{arg}^2 .

has to consider the thresholding mechanism, too. To show that the method is indeed robust against noisy background, we conduct a sensitivity analysis of the threshold. Figure 8¹ shows that if the threshold changes slightly, the number of detected objects does *not* increase or decrease dramatically. Thresholds of 67% to 80% results in two subjects (the parrot and the branch). One of the subjects (the branch) disappears when the threshold rises to 81%. The parrot is detected at the same location for all thresholds from 67% to 100% (i.e., thresholds range of 33%). The branch is detected consistently from 67% to 80% (i.e., thresholds range of 13%). Pay attention to the gradual change of the thresholded responses of D_{arg}^2 in Fig. 8 (which appears as a gradual thinning of the blobs). The theoretic explanation of this robustness is again the fact that Y_{arg} approaches infinity at the negative x -axis of paraboloids; the infinite response there is by far larger than the response obtained by non-convex (or non-concave) image domains.

6. CAMOUFLAGE BREAKING

The previous section described the robustness of the operator under various conditions (illumination, scale, orientation, texture). As a result, the smoothness condition of the detected 3D convex objects can be relaxed. We further increase the robustness demands from the operator by introducing very strong camouflage.

6.1. Biological Evidence for Camouflage Breaking by Convexity Detection

In this section, we exhibit indirect evidence of biological camouflage breaking based on detection of the convexity of the intensity function. This matches our idea of camouflage breaking by direct convexity estimation (using D_{arg}) [20]. We present further evidence, that not only can intensity convexity be used to break camouflage, but also there are animals whose coloring is suited to prevent this specific kind of camouflage breaking.

It is well known that under directional light, a smooth three dimensional convex object produces a convex intensity function. The biological meaning is that when the trunk of an animal (the convex subject) is exposed to top lighting (sun), a viewer sees shades (convex intensity function). As we shall see, these shades may reveal the animal, especially in surroundings which break up shadows (e.g., woods) (see [21]). This biological evidence supports D_{arg} approach of camouflage breaking by detecting the convexity of the intensity function.

The ability to trace an animal based on shadow effects has led, during thousands of years of evolution, to coloration of animals that dissolves the shadow effects. This counter-shading coloration was observed at the beginning of the century [22], and is known as *Thayer's principle of counter-shading*. [21] describes Thayer's principle:

“If we paint a cylinder or sphere in graded tints of gray, the darkest part facing toward the source light, and the lightest away from it, the body's own shade so balances this color scheme that the outlines becomes dissolved. Such graded tints are typical of vertebrates and of many other animals.”

¹Image from: www.freeimages.co.uk







FIG. 8. Threshold sensitivity analysis. **Top Row:** Original image and its D_{arg}^2 map. **Middle Row:** Thresholded D_{arg}^2 maps. **Bottom Row:** Thresholded D_{arg}^2 marked on original images (relative to middle row). Below each image is its threshold value. The complete analysis included threshold values from 67% to 100%; only four of them are displayed.



(a) Constant albedo, top lighting.



(b) Counter-shading albedo, ambient lighting.



(c) Counter-shading albedo, top lighting.

FIG. 9. Thayer's principle of counter-shading. (a) A cylinder of constant albedo under top lighting. (b) A counter-shaded cylinder under ambient lighting. Counter-shading is produced by mapping a convex texture to the cylinder. The texture used here is the convex intensity function: $I(x, y) = (y - 50)^2$, and original texture image is 200×100 pixels. (c) Thayer's principle: the combined effect of counter-shading albedo and top lighting breaks up the shadow effect. In other words, it breaks up the convexity of the intensity function. Animals use this principle to camouflage themselves against convexity based detection.

Figure 9 uses ray tracing to demonstrate Thayer's principle of counter-shading when applied to cylinders. Figure 10 shows some of the animals who use Thayer's principle of counter-shading. The sketches in Fig. 10 demonstrate how animal coloration changes gradually from dark (the upper part) to bright (the lower part). When the animal is under top lighting (sunlight), the gradual change of albedo neutralizes the convexity of the intensity function. Had no counter-shading been used, the intensity function would have been convex (as in Fig. 9(a)), exposing the animal to convexity based detectors (such as D_{arg}). Using counter-shading protects the animal from convexity-based detection.

The existence of counter-measures to convexity based detectors implies that there might exist predators who can use convexity based detectors similar to D_{arg} .

6.2. Thayer's Counter-Shading Against D_{arg} -based Detection

Let us demonstrate how Thayer's principle of counter-shading can be used to camouflage against D_{arg} -based detectors. In Fig. 11 we once again consider a synthetic cylinder; this time we operate D_{arg} on each of the images of that cylinder. As can be seen, the counter-shaded cylinder under top lighting (Fig. 11, Column C) attains much lower D_{arg} values than the smooth cylinder under the same lighting




FIG. 10. Examples of counter-shaded animals (adopted from [21]). The upper part (towards the sun) of the animal is the darker one; transition from the dark part to the bright part is gradual. Here, the animals are under ambient lighting; when the animals come under top lighting (sunlight, in nature), the gradual change of albedo neutralizes the convexity of the intensity function.

(Fig. 11, Column A). This is because counter-shading turns the intensity function from convex to (approximately) planar.

To see the transition from a convex intensity function to a planar one due to camouflage, we draw (Fig. 12-left) the vertical cross-sections of the *intensity* functions of Fig. 11. The smooth cylinder under top lighting (Column A) produces a convex cross-section. The albedo, or the counter-shaded cylinder under ambient lighting (Column B), consists of graded tints of gray, so its intensity function is also convex. Finally, the counter-shaded cylinder under top lighting (Column C) produces a flat intensity function, which means a reduced detection probability by D_{arg} .

Let us verify that the flat intensity function is indeed harder to detect using D_{arg} than the convex intensity function. In other words, we would like to see that under top lighting, D_{arg} 's response to the counter-shaded cylinder is weaker than its response to the smooth cylinder. This can be seen from the vertical cross-sections of the responses of D_{arg} to the various images of the cylinder (Fig. 12-right).

The above demonstrates that Thayer's principle of counter-shading is an effective biological camouflage technique against convexity-based camouflage breakers, and more specifically, against D_{arg} . One can thus speculate that convexity-based camouflage breaking might also exist in nature (or else, camouflage against it would be rendered unnecessary).

The issue of whether or not convexity-based detectors are in fact in use in biological vision systems needs further exploration; this paper presents only indirect evidence for such a usage.

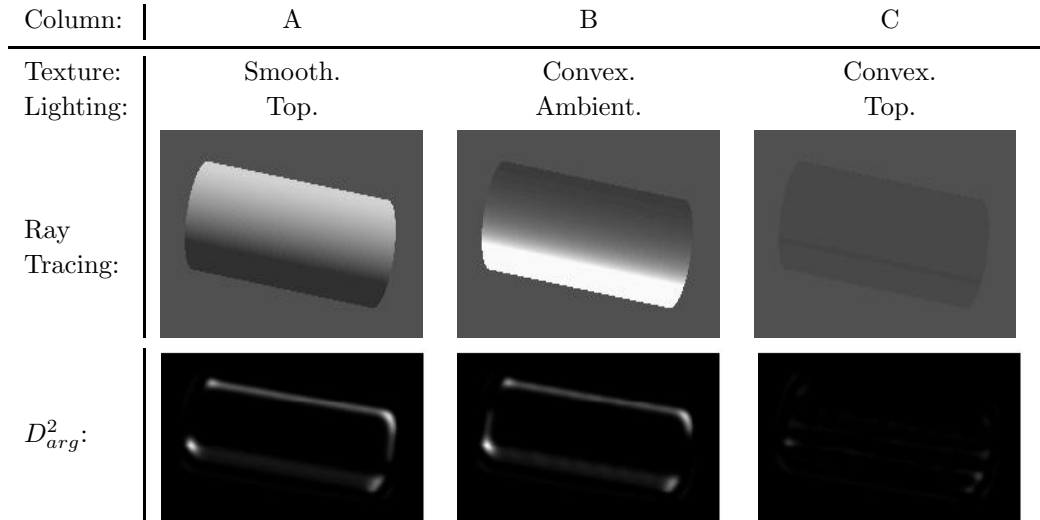


FIG. 11. Operation of D_{arg}^2 on a counter-shaded cylinder. **Column A:** A smooth cylinder under top lighting. **Column B:** The counter-shaded cylinder under ambient lighting. **Column C:** The counter-shaded cylinder under top lighting. The counter-shaded cylinder can barely be noticed under top lighting, due to the camouflage. Under top lighting, the response of D_{arg} is much stronger when the cylinder is smooth than when it is counter-shaded, showing this type of camouflage is effective against D_{arg} .

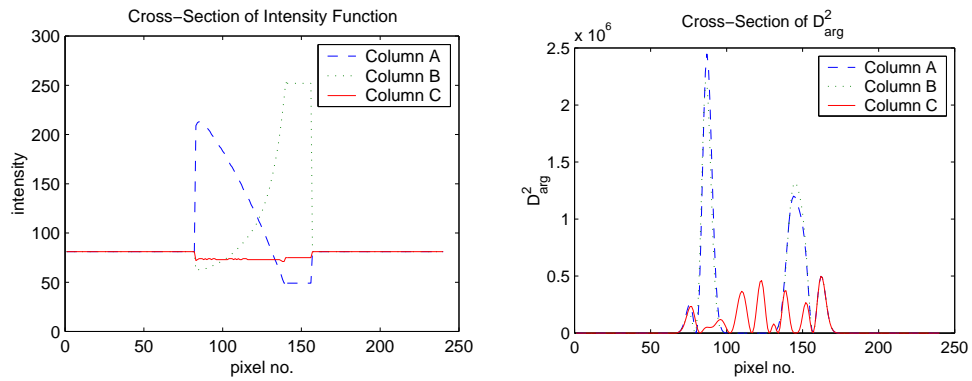


FIG. 12. Cross-sections (parallel to the y -axis, at the center of the image) of: **Left:** The intensity functions of Fig. 11. Thayer's counter-shading yields a flat intensity function for a cylinder. **Right:** D_{arg}^2 of Fig. 11. Under top lighting, the flattened intensity function of the counter-shaded cylinder has a lower D_{arg} response than that of the convex intensity function of the smooth cylinder.

7. CAMOUFLAGE BREAKING DEMONSTRATION

The previous sections have shown both a biological motivation for camouflage breaking by convexity detection, and the high robustness maintained by D_{arg} under harsh conditions (illumination, scale, orientation, texture). The combination of the two makes D_{arg} suitable for usage as a camouflage breaker.

In order to evaluate the performance of the operator in camouflage breaking, we juxtapose convexity based and edge based detection. We chose to compare convexity based detection with edge based detection, because camouflage by edge enhancement is evident in the animal kingdom: [23] (for example) describes camouflage of frogs (*Limnodynastes tasmaniensis*) which superexcites the edge detectors of the vision system of the grater snake (*Thamnophis sirtalis*). This implies that edges are biologically used for camouflage breaking.

In our comparison, D_{arg} is the representative convexity based detector while radial symmetry is the representative edge-based method. The radial symmetry operator searches generalized radial symmetry in the edge map of an image (see [3] for more details). Radial symmetry has been shown to generalize most of the existing methods for interest points detection, including detection of maximal curvature points, corner detection, and edge density measures.

The comparison of D_{arg} with radial symmetry in the task of camouflage breaking deals with several types of camouflage. Each camouflage type is accompanied with a real-life demonstration.


A Technical Note. The radial symmetry operator is scale-dependent, while the peaks of D_{arg} are not. Therefore, we have compared D_{arg} with radial symmetry of radii: 10 and 30 pixels (i.e, 2 radial symmetry transformations performed for each original image). In the paper, only one radius is introduced per original, but similar results were obtained for the other radius as well.

7.1. Indoor Scenes

One of the most basic techniques of camouflage is hiding a single-color object on a background of the same color. Figure 13 shows a white bottle camouflaged on a background of white walls. Some other objects are present in the scene (a computer, a book pile, papers on a desk, etc.). The bottle is covered with a cap, to make it easier for the reader to locate the bottle in the image (the bottle stands on the book pile). The edge-based radial symmetry operator detects other objects (the book pile and the papers), because their edges are much stronger than those of the camouflaged bottle. The edges of the camouflaged bottle are weak and can barely be seen in the edge map. The inherent problem of edge maps of smooth objects is their response to the *outlines* of the smooth 3D object, rather than to the interior of the object (being smooth). Consequently, the surroundings of the object play an important role in determining the strength of the edge. The closer the colors of the subject and the surroundings, the weaker the outline edge. The response of D_{arg} , on the other hand, is to the *intensity surface* formed by the interior of the image of the subject (the bottle, in this case), regardless of the vicinity of the subject in the scene. The distinction that D_{arg} performs between the subject (white bottle) and the background (white wall) is based on an internal feature of the subject:



Original image.



Edge map.

 D_{arg}^2 .Detection by D_{arg} .Radial symmetry ($r=10$).

Detection by Radial Sym.

FIG. 13. A camouflaged white bottle. The white bottle is photographed against a white background. Some other objects are also present in the scene. Edge-based detection schemes disregard the bottle as its edges are very weak (and can hardly be noted in the edge map). D_{arg} , on the other hand, responds strongly to the smooth 3D convex bottle, despite its camouflage.

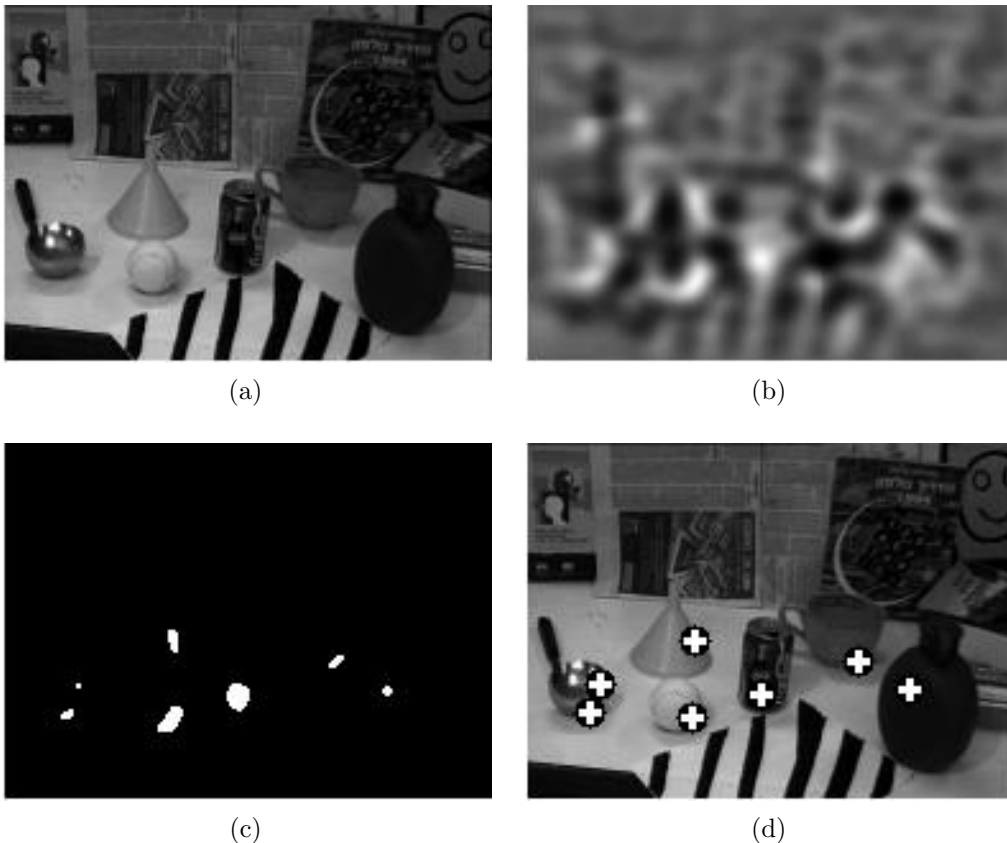



FIG. 14. Attempt to conceal a convex object among other convex objects leads to multiple subject detection. The response of D_{arg} to each convex object is not affected by the presence of other convex objects in the scene. Objects at various scales have been chosen to further demonstrate the robustness. **(a)** A scene containing six smooth 3D convex objects at various scales: ladle, funnel, tennis ball, cola can, mug, and canteen. **(b)** The D_{arg} of (a). **(c)** 70% binary threshold of D_{arg} . **(d)** Detection by 70% threshold of D_{arg} is marked over (a). All convex objects are detected by D_{arg} concurrently.

convexity. Thus, D_{arg} has a strong response to the smooth 3D convex bottle (see D_{arg}^2 map in Fig. 13).


In the above example, the subject assimilates to the environment by mimicking the environment color, which misleads edge-based detection. However, what would happen if one attempts to mislead D_{arg} by mimicking the feature to which D_{arg} responds (convexity)? To test this, we let a convex subject assimilate to a convex environment. As can be seen in Fig. 14, when the scene consists of several convex objects, D_{arg} is able to detect them all. In this example, D_{arg} is not distracted by the presence of other convex objects, as its response is to the internal intensity surface formed by each object. The processing of the local intensity function induced by the object and referred to as a 3D surface along with the use of a robust opera-



a. Rocky Mountain Sheep.

b. D_{arg}^2 .c. Detection by D_{arg} .

d. Edge map.



e. Edge-based detection.

FIG. 15. Rocky mountain sheep in a rocky environment. Apatetic coloration does not distract D_{arg} . (b),(c): D_{arg} detects the sheep, since they are 3D, smooth, and convex. (d),(e): Failure of edge-based methods to detect the sheep: edge-based methods fail, since the edges introduced by the sheep, being convex and using apatetic coloring, are weaker than the edges of the rocky mountain.

tor reduces the environmental effect on the located points, and enables camouflage breaking by D_{arg} .

7.2. Animal Camouflage

Animals use various types of camouflage to hide themselves. One type of camouflage is apatetic coloration. Figure 15 shows two Rocky Mountain sheep (*Ovis*

Canadensis) in their natural rocky environment. The coloration of the Rocky Mountain sheep fits their habitat (pay attention in particular to the upper sheep in Fig. 15(a)). D_{arg} detects the sheep as the main subject (using 70% threshold on D_{arg}^2), since they appear smooth (from the photographic distance), they are three dimensional and convex. Due to the apatetic coloring, the edges of the sheep are weaker than these of the rocky background (Fig. 15(d)), so edge-based methods are doomed to fail on this image. Even the radial symmetry transform — a more sophisticated edge-based transform searching generalized symmetry — fails due to the weakness of the edges of the sheep. As can be seen in Fig. 15(e), radial symmetry detects many background locations. Radial symmetry specifies no single target in the image: 7 different locations were detected by 90% (!) thresholding, spread on a large area of the image.

The habitat of a squirrel demands a more complex camouflage. The squirrel lives in a bushy environment, where colors are not uniform. Leaves, bushes and the ground form a texture which has to be mimicked by the fur of the squirrel. Fig. 16 exhibits a natural camouflage of a squirrel in a leafy environment under the shades of a nearby tree. There are many edges of the squirrel and the environment mixing together, and preventing the radial symmetry operator from isolating any specific target. D_{arg} , however, produces a single strong peak, exactly on the squirrel. The convexity of the squirrel (and in particular, its belly) is the reason for its detection by D_{arg} . The only smooth 3D convex domain in the image is the belly of the squirrel. Though some of the shades might look similar to the belly of the squirrel (even to a human viewer), they do not possess the property of being a projection of a 3D convex object, and thus their intensity function introduce no 3D convexity.

7.3. Military Camouflage



7.3.1. Personnel Camouflage

Spotted uniforms, usually used by military forces, attempt to imitate the distribution of edges and colors in the visual environment of trees, grass, rocks etc. of the arena. This kind of camouflage is exactly the one which deceives edge-based detection. For these cases, edge-based detection is not enough, because the input to the detector is almost identical for background and subject domains.

Although camouflage interferes with the smooth convex intensity function that would otherwise be generated by the subject, in many cases D_{arg} is robust enough to relax the smoothness constraint.

In Fig. 17, a camouflaged hunter is detectable by D_{arg} . The usage of stripes in the uniform does not conceal the convexity of the hunter. Note, that the hunter in Fig. 17 is wearing a very sophisticated disguise which easily misleads human beings (this is the image from the introduction, Fig. 1). Human beings find it difficult to locate the hunter in that image, and nevertheless, D_{arg} detects the hunter (all but a single detection are on the image of the hunter).

In Fig. 18, edge-based methods are able to detect neither the buck nor the hunter. The buck is not detected by these methods, as it is a smooth convex 3D object, and its edges are very weak. The hunter is not detected by edge based methods due to a different reason: he is wearing a camouflage whose colors and edge distribution mimic the colors and edge distribution of the forest. D_{arg} is able to answer both problems simultaneously: The buck is a smooth convex three dimensional object,

Radial symmetry ($r=30$).

Edge-based detection.





 D_{arg}^2 .Detection by D_{arg} .

FIG. 16. A hidden squirrel. The squirrel is on a leafy ground shaded by a tree. The shades and leaves form many edges “deluding” the edge-based methods. Even human viewers find it difficult to locate the squirrel in the image. D_{arg} detects the squirrel, breaking the camouflage.



Original image.

 D_{arg}^2 .Detection by D_{arg} .

Gradient modulus.

Radial symmetry ($r=30$).

Detection by radial symmetry.

FIG. 17. Visual human camouflage. The camouflage mimics the texture and the color of the natural environment. As a result, the domains of the subject and the background woods in the edge map unite and disable edge-based detection. Radial symmetry fails to detect the subject. Convexity based detection distinguishes between the surroundings and subject, since the subject is 3D and convex. Most of the points detected by D_{arg} are of the subject (except for a single outlier).

and is therefore detected by D_{arg} . Taking a threshold of 50% of the maximal D_{arg}^2 detects the buck. If one reduces the threshold, one finds out that at a threshold of 15% of the maximal D_{arg}^2 , the hunter is detected. The hunter is detected since it is also a 3D convex object. Although the convexity of the intensity function of the hunter is disturbed by the presence of camouflage, the hunter convexity is still more dominant in the image than that of the surrounding woods.

Another way to refer to the multiple subject detection of Fig. 18: If we detect the buck and cut its image out of the original image, we remain with an image of a camouflaged hunter. In this image, the hunter can be detected by D_{arg} , but not by edge-based methods (Fig. 19).

Another example of breaking military personnel camouflage appears in Fig. 20. A camouflaged soldier on the background of dense bushes is very hard to detect. Edge based detection misses the soldier. 95%(!) threshold on the output of D_{arg} is necessary in this case to isolate the camouflaged soldier, and still one outlier (out of three locations) occurs.

Another way of concealment is covering most of the body or equipment under branches of trees common in the area of operation. By definition, this ensures the camouflage fits the surroundings. Detecting this camouflage by edge-based means is doomed to fail, as the camouflage consists of many strong edges (tree branches) identical to those of the background. An example for this type of camouflage can be seen in Fig. 21: two soldiers are camouflaged by branches of trees. Edge-based techniques detect background domains. D_{arg} detects the soldiers, as their (uncovered) faces are 3D and convex. See [24, 25] for details on face detection using D_{arg} .

7.3.2. Military Camouflage

Camouflage and concealment of equipment is of particular interest. This section continues the comparison of detection by D_{arg} and by radial symmetry for camouflaged artifacts.


Figure 22 present a tank in camouflage paints in front of a tree. The tree produces edges distracting the radial symmetry operator. The convexity of the intensity function near the wheels of the tank exposes the tank to the D_{arg} detector.

Figure 23 presents tanks in a highly cluttered scene. The edges of the urban area distract edge-based detectors. The convexity of the tanks leads to detection of 2 out of 3 tanks by D_{arg} (the third tank is too small to detect).

8. CONCLUSIONS

We have defined a convexity-based operator, D_{arg} , for detection of image domains emanating from smooth *convex* (or concave) three dimensional objects. D_{arg} is not based on edge maps, and is thus free of their flaws (e.g., it is robust in dominant textures and camouflage). D_{arg} is proved invariant under any derivable strongly monotonically increasing transformation of the image intensity function, which practically means robustness to illumination changes. Robustness to orientation and scale is also explained. We have shown that D_{arg} is suitable for camouflage breaking, and that in practice the smoothness assumption can be relaxed.


Having D_{arg} defined, we have illustrated how Thayer's principle of counter-shading prevents detection based on the convexity of the intensity function of the



a. Original image. b. D_{arg}^2 . c. Detection by D_{arg} (50% thresh).



d. 15% threshold of D_{arg}^2 . e. Detection by D_{arg} (15% thresh). f. Edge map.



g. Radial symmetry (r=10). h. Detection by 90% thresh of Radial Sym. i. Detection by 80% thresh of Radial Sym.


FIG. 18. A camouflaged hunter and a buck in the woods. D_{arg} detects both the buck (by 50% of D_{arg}^2) and the hunter (by 15% of D_{arg}^2).



Original image.



Edge map.

 D_{arg}^2 .Detection by D_{arg} .

Radial symmetry (r=10).



Detection by Radial Sym.


FIG. 19. A camouflaged hunter. The buck's image was cut out of Fig. 18, so that only the camouflaged hunter and the background remains. The natural environment and the hunter's camouflage unite and form a uniform-looking edge map. Detection techniques which use this edge map do not receive enough information to be able to distinguish the subject from the background. D_{arg} detects the subject based on convexity information: the hunter is a convex 3D object.



a. Original image.




d. Edge map.



b. D_{arg}^2 .



c. Detection by D_{arg} (95% thresh).



e. Radial symmetry ($r=10$).



f. Detection by 95% thresh of Radial Sym.


FIG. 20. A camouflaged soldier on the bank of a river near dense bushes. Edge distribution unites with camouflage (especially the helmet). In spite of the strong camouflage, D_{arg} detects the soldier in 2 out of 3 locations it isolates. The incorrect detection is of clouds reflecting in the water.



Original image.




Edge map.


 D_{arg}^2 .Detection by D_{arg} .Radial symmetry ($r=10$).

Detection by Radial Sym.

FIG. 21. Camouflaged soldiers. The camouflage trees produce many strong edges, disguising edge-based detection schemes from the soldiers. 80% thresholding of the radial symmetry map does not detect any specific region of interest. Nevertheless, the soldiers faces are 3D and convex, and thus lead to their detection by D_{arg} .



Original image.



Edge map.





 D^2_{arg} .Detection by D_{arg} .

FIG. 22. A tank in camouflage paints near a tree. Convexity-based D_{arg} is *not* distracted by the background (tree, grass).

Radial symmetry ($r=10$).


Detection by Radial Sym.

Radial symmetry ($r=30$).

Detection by Radial Sym.

FIG. 22—*Continued*

A tank in camouflage paint near a tree. The tree distracts edge-based detection, regardless of the scale.

Radial symmetry ($r=10$).

Detection by Radial Sym.


 D_{arg}^2 .Detection by D_{arg} .

FIG. 23. Tanks in an urban zone (i.e., high clutter). Edge-based detection is distracted by the background. Convexity-based D_{arg} detects 2 of the 3 tanks.

subject (using D_{arg}). This gives further support to the speculative conclusion, that D_{arg} might be used in a biological vision system of a predator whose detection ability is based on convexity detection, and the apatetic coloration of its prey is based on Thayer's principle of counter-shading.

Camouflage breaking by D_{arg} has been demonstrated to be highly effective. This has been shown in a comparison between a convexity-based camouflage breaker (D_{arg}) and a representative edge-based operator (radial symmetry). The comparison included various types of camouflage: indoor scenes, animal and human camouflage and military camouflage. Convexity-based camouflage breaking has been found to be highly robust and in most of the cases much more effective than edge-based techniques.

Despite the better results obtained by D_{arg} in the comparison, one must bare in mind that there still exists a camouflage technique which overcomes D_{arg} : Thayer's principle of counter shading. Consequently, D_{arg} is suggested as a *complementary* method to edge-based techniques, so that each discipline benefits from the advantages of the other.

REFERENCES

1. A. C. Copeland and M. M. Trivedi. Models and metrics for signature strength evaluation of camouflaged targets. *Proceedings of the SPIE*, 3070:194–199, 1997.
2. S. W. Zhang and M. V. Srinivasan. Prior experience enhances pattern discrimination in insect vision. *Nature*, 368(6469):330–332, March 1994.
3. D. Reisfeld, H. Wolfson, and Y. Yeshurun. Context free attentional operators: the generalized symmetry transform. *International Journal of Computer Vision*, pages 119–130, 1995.
4. F. M. Gretzmacher, G. S. Ruppert, and S. Nyberg. Camouflage assessment considering human perception data. *Proceedings of the SPIE*, 3375:58–67, 1998.
5. L. B. Hepfinger. Camouflage simulation and effectiveness assessment for the individual soldier. *Proceedings of SPIE*, 1311:277–283, 1990.
6. G. E. Irvin and M. G. Dowler. Physiologically based computational approach to camouflage and masking patterns. *Proceedings of SPIE*, 1700:481–488, 1992.
7. D. H. Kelly. 40 years of image technology and vision research. *Optics and Photonics News*, pages 913–922, July 1990.
8. S. P. McKee, S. N. J. Watamaniuk, J. M. Harris, H. S. Smallman, and D. G. Taylor. Is stereopsis effective in breaking camouflage? *Vision Research*, 37:2047–2055, 1997.
9. I. V. Ternovskiy and T. Jansson. Mapping-singularities-based motion estimation. *Proceedings of the SPIE*, 3173:317–321, 1997.
10. L. Huimin, W. Xiuchun, L. Shouzhong, S. Meide, and G. Aike. The possible mechanisms underlying visual anti-camouflage: a model and its real-time simulation. *IEEE Transactions on Systems, Man & Cybernetics, Part A (Systems & Humans)*, 29(3):314–318, May 1999.
11. S. Marouani, A. Huertas, and G. Medioni. Model-based aircraft recognition in perspective aerial imagery. In *Proceedings of the International Symposium on Computer Vision*, pages 371–376, Coral Gables, Florida, USA, November 1995.
12. S. Guilan and T. Shunqing. Method for spectral pattern recognition of color camouflage. *Optical Engineering*, 36(6):1779–1781, June 1997.
13. T. W. Cronin, N. Shashar, and L. Wolff. Portable imaging polarimeters. In *Proceedings of the 12th IAPR International Conference on Pattern Recognition*, pages 606–609, Jerusalem, Israel, October 1994.
14. N. Shashar, R. Hagan, J. G. Boal, and R. T. Hanlon. Cuttlefish use polarization sensitivity in predation on silvery fish. *Vision Research*, 40(1):71–75, 2000.
15. S. W. Zucker, M. S. Langer, L. A. Iverson, and P. Breton. Shading flows and scenel bundles: A new approach to shape from shading. In G. Sandini, editor, *Second European Conference on Computer Vision '92*, Santa Margherita Ligure, Italy, May 1992. Springer-Verlag.
16. D. H. Ballard and C. M. Brown. *Computer Vision*, pages 123–131. Prentice-Hall Inc., New Jersey, USA, 1982.
17. J. C. Russ. *The Image Processing Handbook*, pages 242–244. CRC Press, 2 edition, 1995.

18. J. Bigün. Recognition of local symmetries in gray value images by harmonic functions. In *The 9th International Conference on Pattern Recognition*, pages 345–347, Rome, Italy, November 1988. IEEE Computer Society.
19. A. Tankus and Y. Yeshurun. Detection of regions of interest and camouflage breaking by direct convexity estimation. In *IEEE Workshop on Visual Surveillance*, pages 42–48, Bombay, India, January 1998.
20. A. Tankus and Y. Yeshurun. Convexity-based camouflage breaking. In *15th International Conference on Pattern Recognition*, Barcelona, Spain, September 2000. To appear.
21. A. Portmann. *Animal Camouflage*, pages 30–35. The University of Michigan Press, 1959.
22. A. H. Thayer. An arraignment of the theories of mimicry and warning colours. *Popular Science Monthly, N. Y.*, pages 550–570, 1909.
23. D. Osorio and M. V. Srinivasan. Camouflage by edge enhancement in animal coloration patterns and its implications for visual mechanisms. *Proceedings of the Royal Society of London B*, 244:81–85, 1991.
24. A. Tankus, Y. Yeshurun, and N. Intrator. Face detection by direct convexity estimation. In Josef Bigün, Gérard Chollet, and Gunilla Borgefors, editors, *Proc. 1st Intl. Conf. on Audio- and Video-based Biometric Person Authentication*, pages 43–50, Crans-Montana, Switzerland, March 1997. Springer.
25. A. Tankus, Y. Yeshurun, and N. Intrator. Face detection by direct convexity estimation. *Pattern Recognition Letters*, 18:913–922, 1997.