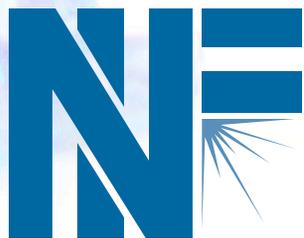


Efficient Detection of Objects of Unknown Size and Shape in Noisy Images

Marijn Bezuijen
Laura Kegelmeyer
Judy Liebman

NIF Optics Inspection
Analysis Team



UCRL-PRES-208126

Background

- Minor imperfections in NIF optics can scatter laser light
- Imperfections that are detected early can be mitigated to prevent them from causing damage

Example of an imperfection seen by SIDE



Astronomy Data Will Be Used in this presentation

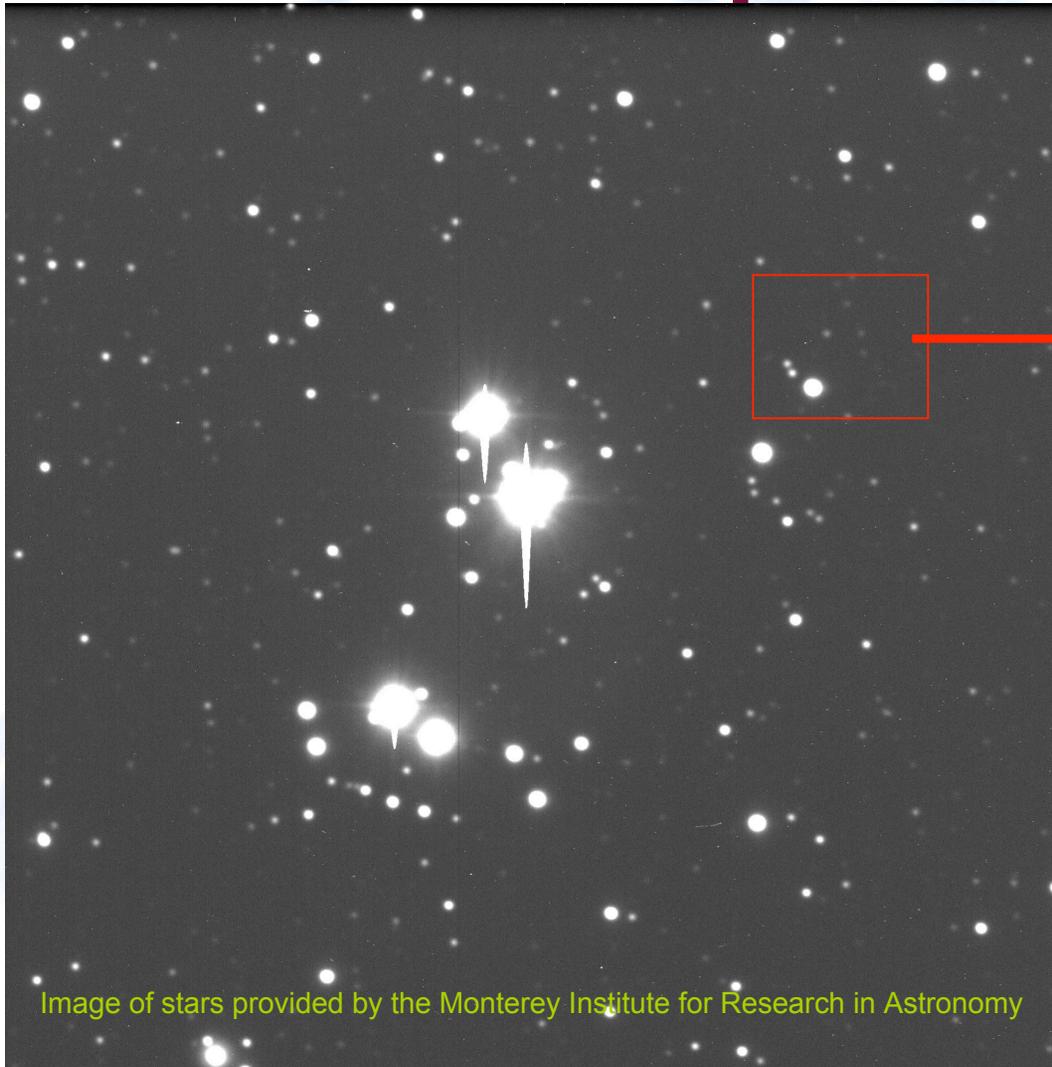
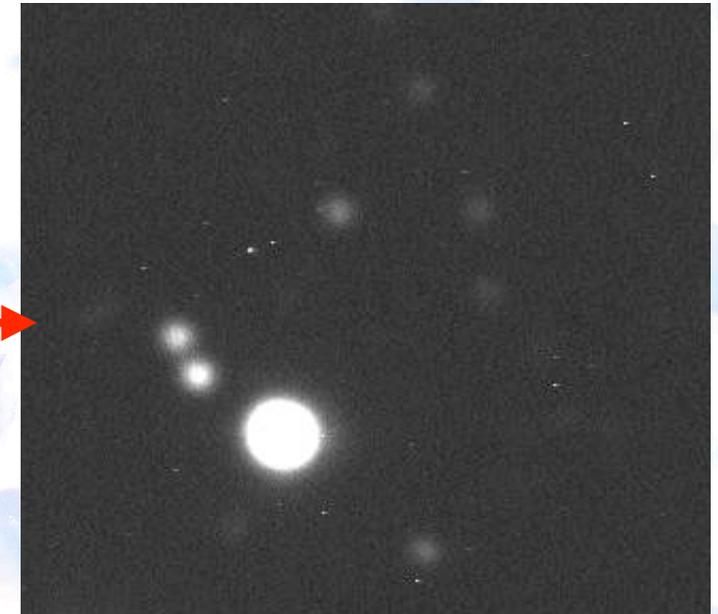


Image of stars provided by the Monterey Institute for Research in Astronomy



Contains many objects of different sizes and contrast levels combined in one noisy image, provides far more visible examples.

An Ideal Detection Algorithm

- Finds features of all sizes
- Ignores background, ghosts and noise
- Is independent of image properties (such as size, bit depth, orientation)
- Is insensitive to changes elsewhere in the image
- Is comprehensible and predictable
- Has a simple and fast implementation

What's needed

- A method that will separate object pixels from background, noise and ghosts, independent of any image properties

Many common segmentation methods don't work

- Threshold
 - Impossible to choose
- Background subtraction
 - Scale biased
- Template match
 - Many shapes possible
- Linear filters
 - Frequencies unknown
- Morphological filters
 - Scales unknown

Nonetheless, the human brain can do it ..

So how does our brain do it ?

- It looks for things that 'stand out' from both background and noise



Local signal to noise ratio is the key

Signal to Noise Ratio, the Basics

- We define signal to noise ratio for each object as:

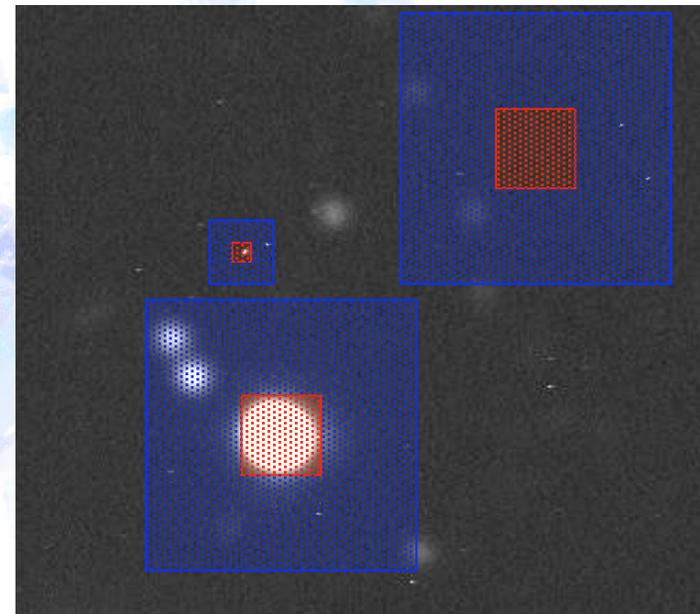
$$SNR = \frac{\mu_s \sigma_{bg}}{\sigma_{bg}} \quad (1)$$

Where: m_s is the mean signal value inside the object
 m_{bg} is the mean of local background around the object
 s_{bg} is the standard deviation of the background around the object

- However, we don't know yet where our objects and background are in the image, so we need a suitable estimator of the SNR for all pixels in the image

Local SNR Can Be Estimated (Current Method)

- Look at each pixel in the image
- For that pixel, estimate the local SNR
- If the small neighborhood contains an object, the SNR will be much higher
- The resulting image can be thresholded depending on requirements
- Different window sizes can be used to find differently sized defects
- This method has been successfully used in NIF



Local SNR can be estimated using full image logic

- We define the following estimator for the SNR for each pixel:

$$SNR(i, j) = \frac{S_{bgs}(i, j)}{S_{\square}(i, j)} \quad (2)$$

Where
$$S_{bgs}(i, j) = S(i, j) - \sum_{m,n} c_{mn} S(i+m, j+n) \quad (3)$$

= a background subtracted version of $S(i, j)$

And
$$S_{\square}(i, j) = \sqrt{\sum_{m,n} k_{mn} S_{bgs}(i+m, j+n)^2} \quad (4)$$

= the estimated local standard deviation in $S_{bgs}(i, j)$

And
$$\sum_{m,n} c_{mn} = \sum_{m,n} k_{mn} = 1 \quad (5)$$

Define the size and shape of the windows for background subtraction and standard deviation

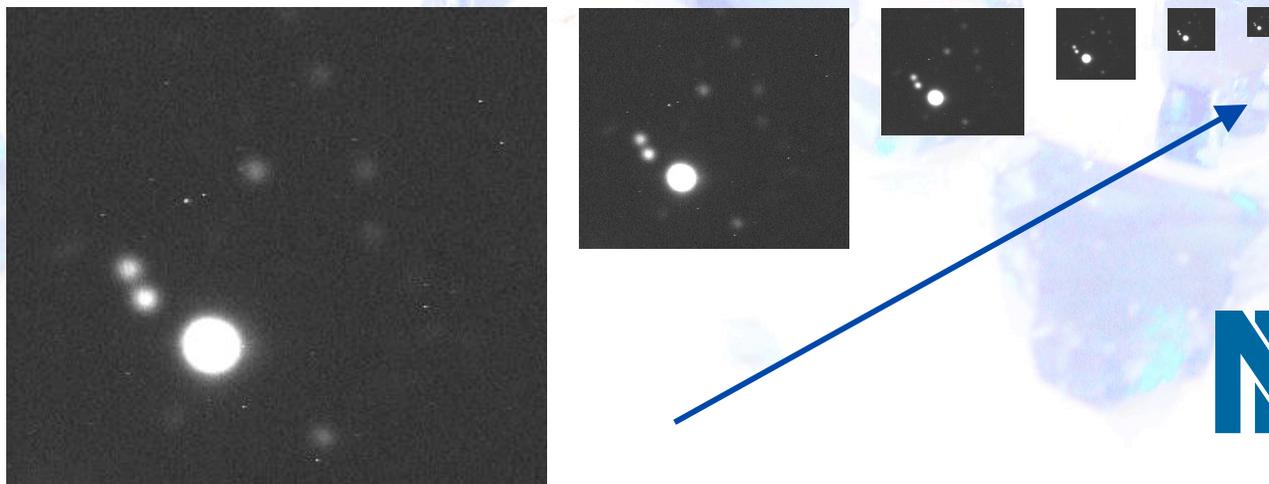
Choosing the weighting factors

- For a proper estimation of the standard deviation, the window used for background subtraction should be smaller than the window used for estimation of standard deviation
- A circularly symmetric, smooth window can be obtained by using Gaussian weighting factors
- A Gaussian averaging filter can be employed to do this, which is linearly separable, resulting in fast implementations
- Weighting the center is desirable as this emphasizes disturbances closer to the object

Extending this algorithm to work on multiple scales

The following recipe can be used to find objects of all scales present in the image:

1. Run the algorithm with a very tight background subtraction (small window), removing all but the smallest objects
2. Scale the original image down by a factor 2
3. Re-run the algorithm at the smaller scale



Scaling down the image has many practical advantages

- Only a single set of parameters is needed to find objects at all scales
- Since detection speed is quadratic with respect to image size, smaller images don't add much time
- If high detection certainty is required, intermediate steps can be done by smoothing the images with a small kernel and re-running the detection with slightly bigger filters

A practical implementation

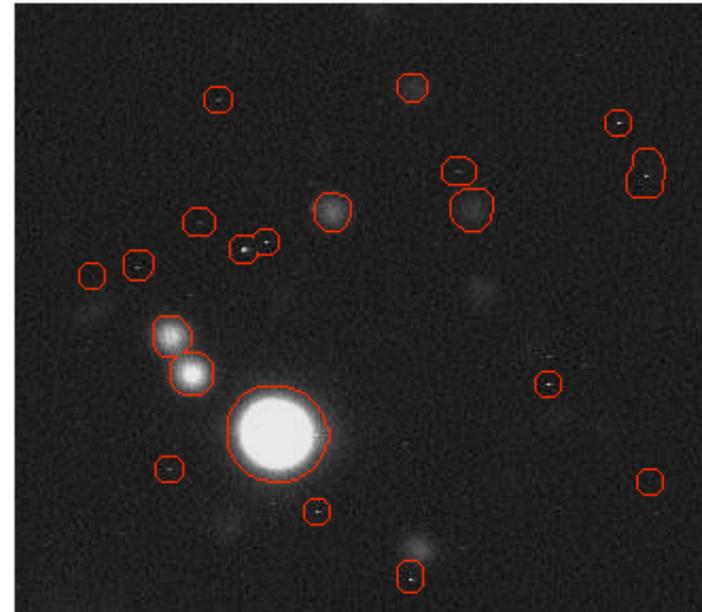
- Find all objects at different scales
- Scale smaller scale detections up to larger scale
- Combine results using a logical OR

After the locations of objects in the image are known:

- Find extent of found objects
- Measure properties of found objects (peak value, area, summed intensity etc.)

The algorithm performs well in practice

- Objects of all sizes are found reliably in many different kinds of input images
- Implementation in MATLAB code
- Runs faster than other tested algorithms



Sample result obtained using full image logic at multiple scales

The Algorithm Is Not Perfect, But Outperforms Competition

Two imperfections we have found are:

- Extent of objects is unknown after detection, a separate step is required to get this information

Separation allows for better tuning of the two steps

- Small, dim objects right next to sharply defined, bright big objects are not always detected, as they are assigned a lower SNR than if there were no big object nearby

In inspection of objects, this is hardly a problem

Conclusion

- We have a reliable and predictable algorithm to find objects in an image that matches how we see objects and can meet NIF detection requirements
- This algorithm can be implemented through only a limited number of operations on a full image, allowing for inspection of all optics on all beamlines