Eliminating Illusions in Microscopy
Basic of Pattern Recognition
By Matthew Putman

Preface

The human mind has abilities which far exceed anything that modern computation can achieve. This is most evident in pattern recognition, which includes most sensory inputs. It is only humans and other animals that can truly fill in missing sensory information to such a strong degree as to recognize characters, sounds, faces and more in nearly all locations and times. Much of modern machine learning has been dedicated to computationally replicating the power of the brain in order to eventually exceed it. Ray Kurzweil details some current approaches to the problem in his 2012 book "How to Create a Mind". In the book he discusses the now applied use of big data and statistical modeling to provide recognition of sensory input. This is an attempt to gain the necessary information to mimic the cortical hierarchy that makes humans so powerful in this regard. The success of this type of approach, which uses probability theory such as Bayesian logic and Hidden Markov Models, can be seen in everyday products that are made by Google and others. This is especially apparent in the facial recognition Google uses with Picassa and in Google Translate. Both of these rely on the very large quantities of data generated by users to gain a human-like memory. An even more impressive example of machine success is the IBM computer Watson, which successfully used learning techniques including the above mentioned modeling, and powerful processing to defeat the world's best Jeopardy players.

Despite these exponential improvements in computing power, there are areas that remain a challenge in computational image recognition. These can be thought of in two categories, both of which I will address here:

1. Humans are imperfect; therefore brain mimicry can be imperfect.
2. Most successful models for image recognition involve very large sample sets.

Before addressing this, I will narrow the scope of the discussion. The work of creating any significant general Artificial Intelligence involves many approaches, including the powerful mechanisms described above. There are specific tasks that need to be fulfilled which exceed the human brain's capability. Many times the need to recognize features does not come with large data sets available for solving the generalized problem of machines such as the types of data that Watson had access to. In particular I will focus on a very old, but still ubiquitous area of recognition of optical microscope features. Whether in the field of inspection or pathology, the human eye is often considered the standard for visual recognition. I will present approaches that suggest that now algorithms eliminate some key illusory problems in perception, and therefore create more reproducible microscope imaging systems.

Introduction

Optical illusions are both well documented and well understood. Retinal cortical processing, while efficient, produces some inaccuracies, even where simple edge detection is concerned. One of the most famous examples is the Müller-Lyer Illusion in which the lines of three arrows appear different, despite being the same length.

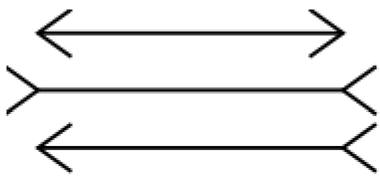

Experience cannot train an observer to overcome even this simple optical illusion. It also does not take computation to discern that the lines are the same size, as the visual illusion can be broken by creating a visual counterpart by making the lines a different color than the ends.

As geometries and morphologies become more complex, this alignment can be used through simple algorithms for clarification and classification. Not only will the illusion no longer be relevant, but numerical description possible. An example where a tool is required for understanding an illusion can be seen below.

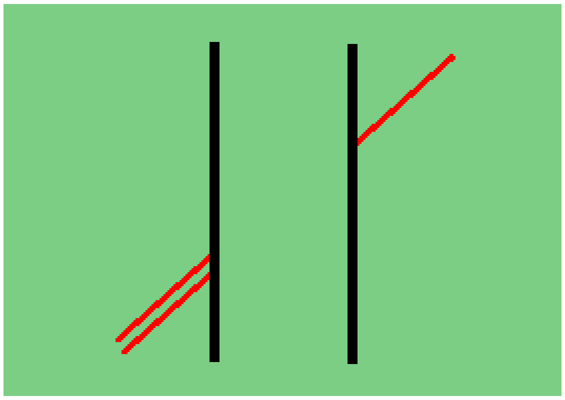

The question of which red line on the left aligns with the red line on the right is illusory. A basic tool for solving the question of alignment is a ruler. In real world applications however two obvious problems exist:

1. An observer does not know what illusion to quantify without being told. That is, an observer does not know whether what they are seeing is an illusion, as nature presents both real and illusory objects, just as artists and geometers do.
2. The quantities of forms in any given sample of material (natural or synthetic) provide too much information for the observer to reasonably work with.

The history of pattern recognition goes back nearly 40 years, however modern approaches have relied heavily on more recent advances such as hyper threading, neuro networks, and large storage requirements.

In the examples above, from a computational perspective, pattern recognition can be thought of as a series of conditions, usually known as input nodes. Input nodes produce desired quantitatively accurate results, often called output nodes. Pattern recognition is generally categorized according to the type of learning procedure used to generate the value of the output nodes. Three very general types of pattern recognition are considered:

1. Artificial learning approach – This method requires a large amount of reference examples. An input node is created by statistically characterizing as many examples of the image as possible. This is known as training data. A learning procedure informs a model that attempts to meet two objectives: Perform as well as possible on the training data, and generalize as well information for new data. Prediction is a challenge in any science, and generally requires a combination of the correct model, and an appropriately large sample set, in this case training data.
2. Classification and clustering – This form of pattern recognition is used when relatively small training sets are available. While initial conditions can be a challenge in classification, small segments of code are written in order to define a desired output. In other words, the numerical descriptions become the inputs, rather than a training set. This is sometime referred to as "hand-labeling" despite the fact that the characteristics of the desired pattern output are determined through automated measurement.
3. Hybrid recognition – this is not a standard term; however the concept is used in nearly all complex pattern recognition techniques. The idea is that the classifications derived in the second approach are used as the training set for the first approach. It is also possible in a realistic setting that certain features can be trained using learning algorithms, while others having a smaller sample set cannot. Therefore a combination of the two is included in the algorithm.

Most common pattern recognition algorithms use variations on probability theory that ranges from Bayesian Models to new methods of modeling such as Agent Based models. This would mean an output node that has an inherent probability associated with it. Unlike other algorithms, which simply output annotated results, pattern recognition tools often output a probability, which can be used as a process control gate in the case of material characterization. When #3 described above is utilized a system can be trained by user labeling, which then outputs a set of possible classifications, rather than a numerical probability.

Microscopy and Pattern Applications

In microscopy humans are often over confidant with their own expertise and ability to recognize features. Using the strategies above, various types of compound semiconductor features were classified. Though image recognition is possible with anything, as a computer does not inherently care what it is looking at, compound semiconductors present several interesting challenges. Here are a few examples of classification challenges that would impossible for the unaided human.

EXAMPLE 1

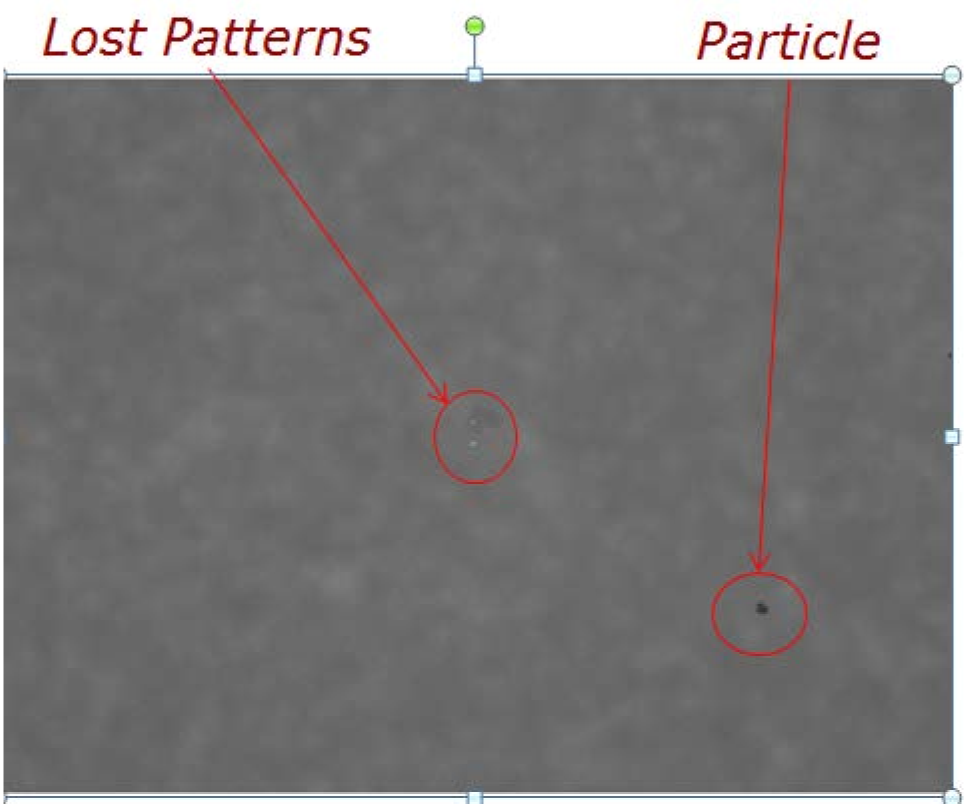

The above image is an optical high resolution image of an area of a patterned Sapphire wafer at 5X magnification. This image is one of over 2000 images which are stitched together in order to have a complete composite of the wafer. The relatively low 5X magnification is therefore preferable for providing a rapid automated image of the entire wafer, as higher magnification would require a greater number of images, and longer imaging time. The two circled areas represent two entirely different features, one of which is a surface particle (external dust), and the other representing two missing parts the pattern. It was confirmed that the image recognition algorithm was accurate by looking at the same area at 20X magnification.

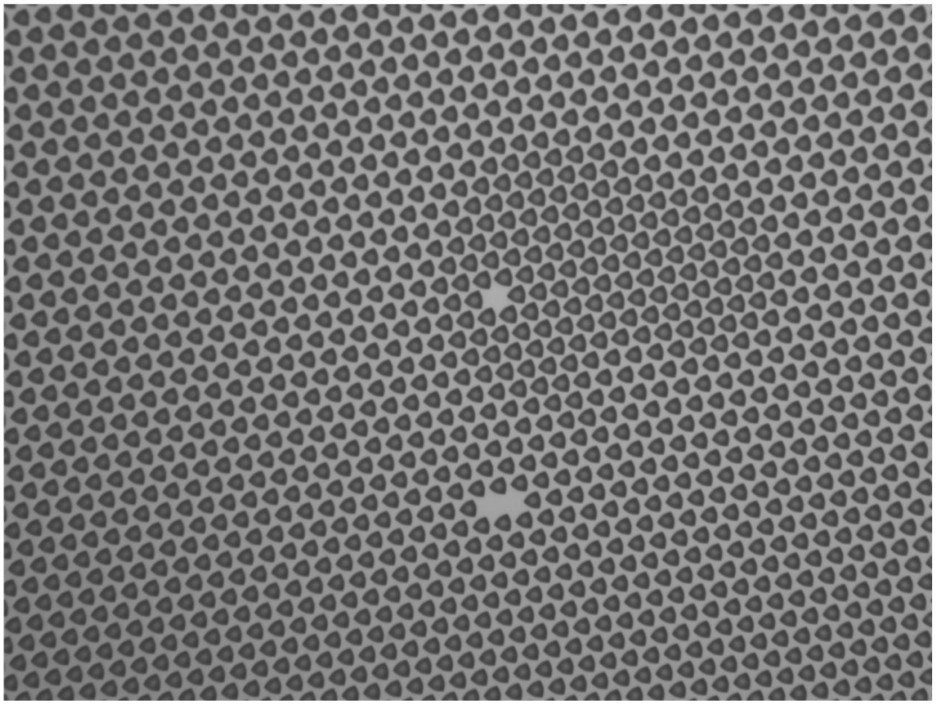

EXAMPLE 2

The following image is another single full resolution 5X image of the surface of a compound semiconductor. The features to the eye all look similar, yet the image recognition classified a particular crystal defect that the user wanted to identify. Those that were not circled, while similar looking, were algorithmically and later through other testing methods, determined to be different features. The algorithm was therefore successful.

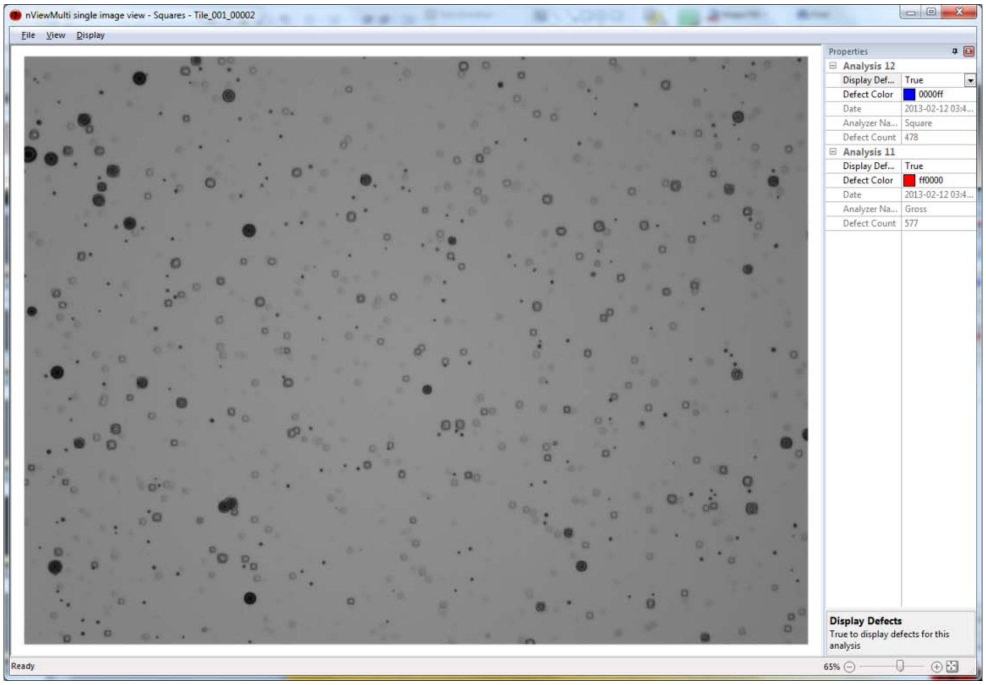

The following is a post image recognition annotation of the features of interest.

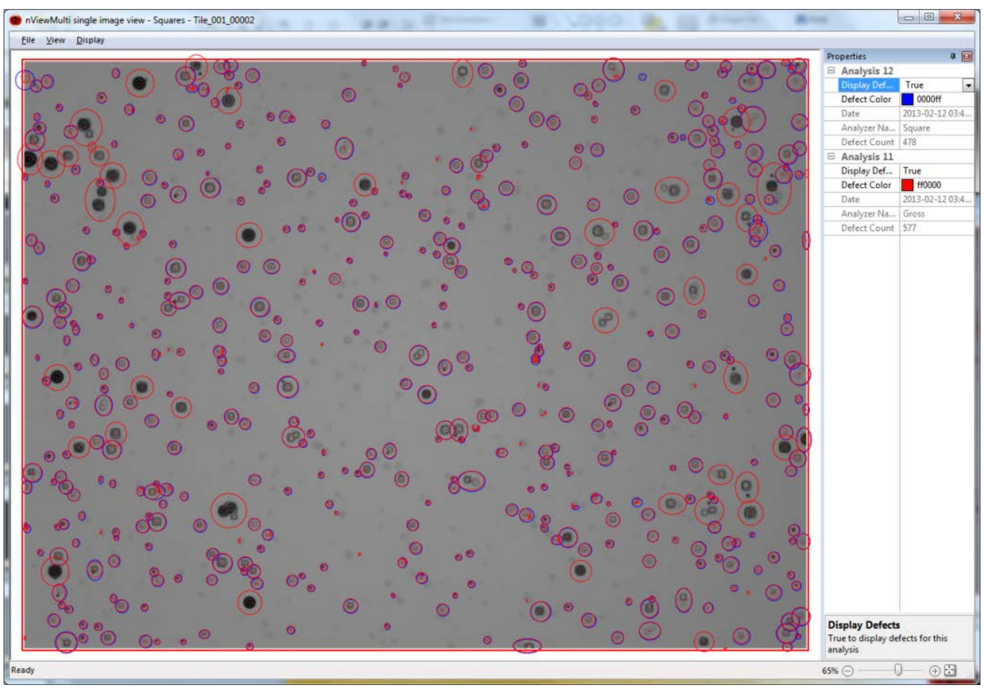

As a test of just how inaccurate human perception can be, I have included this same image, with the same classification, but in three alternative colored backgrounds. The background at which humans view an image changes our perception, where it has no effect on the imagine recognition output.

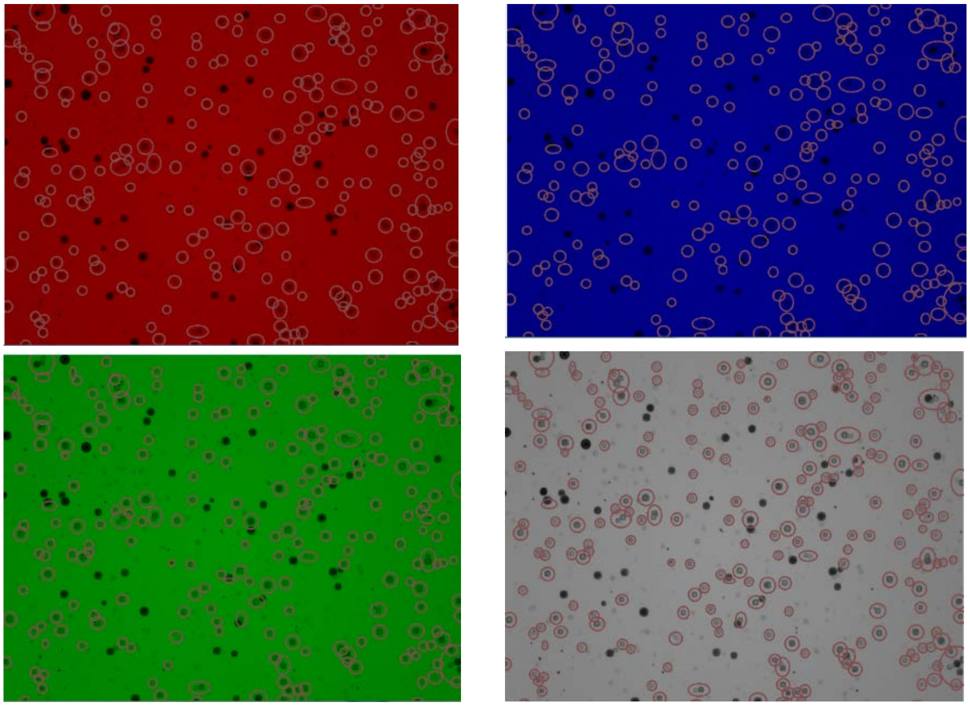

Conclusion

It is widely understood that humans and other anaimals have good visual memory. It is suggested here however that visual memory is only a part of the system required for advance image recognition. A combintion of computational approaches are required to analyzes small features, which have a high degree of similarty.