

Addressing supply chain risks of microelectronic devices through computer vision

Zhenhua Chen¹, Tingyi Wanyan¹, Ramya Rao¹, Benjamin Cutilli¹, James Sowinski¹,
David Crandall¹, and Robert Templeman²

¹School of Informatics, Computing, and Engineering, Indiana University, Bloomington IN
Email: {chen478, tiwanyan, ramrao, bvcutill, japasowi, djcran}@indiana.edu

²Naval Surface Warfare Center, Crane Division, Crane, IN
Email: robert.templeman@navy.mil

Abstract—Microelectronics are at the heart of nearly all modern devices, ranging from small embedded integrated circuits (ICs) inside household products to complex microprocessors that power critical infrastructure systems. Devices often consist of numerous ICs from a variety of different manufacturers and procured through different vendors, all of whom may be trusted to varying degrees. Ensuring the quality, safety, and security of these components is a critical challenge. One possible solution is to use automated imaging techniques to check devices' physical appearance against known reference models in order to detect counterfeit or malicious components. This analysis can be performed at both a macro level (i.e., ensuring that the packaging of the IC appears legitimate and undamaged) and a micro level (i.e., comparing microscopic, transistor-level imagery of the circuit itself to detect suspicious deviations from a reference model). This latter analysis in particular is very challenging, considering that modern devices can contain billions of transistors. In this paper, we review the problem of microelectronics counterfeiting, discuss the potential application of computer vision to microelectronics inspection, present initial results, and recommend directions for future work.

I. INTRODUCTION

Integrated circuits pervade nearly every electronic device, from small kitchen appliances, to consumer smartphones and laptops, to the systems that power critical national infrastructure. Ensuring the security of these electronic devices is of critical importance to government and industry alike. These devices' safe and reliable operation depends on the quality, safety, and security of integrated circuits, but these devices are often made by third-party companies and navigate a complex and insecure supply chain.

A key vulnerability is that defective, counterfeit, or malicious electronic components could be introduced along the way. This is not just a theoretical possibility; instances have in fact become surprisingly common [7]. For example, through a survey of major companies and government organizations, the U.S. Department of Commerce's Bureau of Industry Security concluded that "an increasing number of counterfeit incidents [is] being detected, rising from 3,868 incidents in 2005 to 9,356 incidents in 2008" [16]. Meanwhile, the U.S. Senate Armed Services Committee "uncovered approximately 1,800 cases of suspect counterfeit electronic

parts being identified by some companies in the defense supply chain, with the total number of suspect parts exceeding 1 million" [1]. Perhaps most alarming, the Semiconductor Industry Association has estimated that "as many as 15% of all spare and replacement semiconductors purchased by the Pentagon are counterfeit," and that "counterfeiting costs U.S.-based semiconductor companies more than \$7.5 billion per year, which translates into nearly 11,000 lost American jobs" [2].

Compounding the difficulty of detecting defective and counterfeit parts is the enormous complexity of modern devices, which can contain tens of billions of transistors in a single package [25]. One of the most successful approaches for ensuring that an IC is legitimate and made to proper specifications is visual inspection [38]. While simply inspecting exterior packaging may be useful for detecting obvious quality or legitimacy problems, more careful inspection requires imaging the silicon of the IC at an atomic scale. State-of-the-art imaging techniques like scanning electron microscopes (SEMs), tunneling electron microscopes (TEMs), scanning optical microscopes (SOMs), and microtomography (micro-CT) allow such transistor-level images to be collected, but actually analyzing the imagery produced by these techniques must be performed by human operators, which is extremely costly and labor intensive.

Given the success of computational image analysis techniques in other domains, it is natural to apply the latest computer vision techniques to analyzing microelectronic imagery. In this paper, we discuss the potential application of modern computer vision techniques to the microelectronic inspection task. We first give background on microelectronic counterfeiting, including motivations and typical signals of counterfeit devices. We then review techniques for detecting counterfeits, with a particular focus on visual inspection. This inspection could be performed on the external properties of an IC's packaging, or on its internal circuits using microscopy imaging; we review the current state-of-the-art in imaging techniques that are applicable to the latter. As a first step towards investigating computer vision for IC counterfeit detection, we consider the problem of unsupervised clustering of ICs based on similarity of their external packaging. We present

a deep machine learning-based technique for fine-grained visual comparison, and apply it to a dataset of annotated, real-world IC images that we have collected and have made publicly available. Finally, we recommend fruitful areas for further research.

II. COUNTERFEIT MICROELECTRONICS

Almost all modern electronic devices, from the simplest to the most complex, contain at least one integrated circuit (IC). Considering their critical role in our everyday lives, it is easy to forget that ICs are a relatively young technology, invented less than 60 years ago [3]. In this period of time, the IC has advanced from single transistor circuits in 1958 to complex dies that contain *billions* of transistors today. As ICs have become more advanced, they have grown exceedingly complex, packing in more transistors while shrinking their geometries to the limits of physics. In fact, the term *microelectronics* understates the sophistication of modern devices: today’s ICs could accurately be described as ‘*angstromelectronics*,’ since even a single atom transistor has been successfully demonstrated [4]! Of course, such complex devices are extremely complicated to manufacture.

A. Reliability and integrity of microelectronics

Despite improvements in the reliability of modern microelectronics, failures still happen and can have catastrophic consequences. Some of these failures are caused by inevitable variations in quality during manufacturing. But compounding the problem is that modern devices typically consist of multiple ICs, often manufactured by different companies in different parts of the world, that traverse complex supply chains during the manufacturing and delivery process. These complex supply chains promote efficiency of scale and cost effectiveness, but they also introduce many opportunities for devices to be intercepted, modified, or replaced [40].

The causes and motivations of these device substitutions can be complex – see [38] for a thorough discussion – but generally fall into three broad categories. Some substitutions may be inadvertent: a well-intentioned manufacturer may accidentally ship the wrong part, for example, and the error may remain undetected throughout the supply chain. Financial motivations are another major cause. An opportunistic supplier may substitute a less expensive part, for example, or may provide recycled or rejected parts that are represented as new. The sophistication of these counterfeiting techniques varies considerably, but are often quite simple, like sanding off or painting over an inexpensive IC’s markings and replacing them with the part number of a more expensive device.

A third (and more alarming) motivation is for a malicious adversary to purposely inject custom ICs into the supply chain, in order to change the functionality of the devices that use them. These custom ICs may be designed to fail early, to maliciously alter the functionality of a device, to allow the device to be controlled remotely, or even to surreptitiously steal information. In one particularly striking example, Yang et al. [43] demonstrated an attack in which a few extra gates

were added to the die of a microprocessor. This extra circuitry had no apparent effect on the device’s functionality, and could easily go unnoticed amongst the billions of transistors on a modern integrated circuit. However, when a software trojan executed a particular series of instructions, the extra gates would give the trojan privileged access to the microprocessor, allowing it unfettered access to all data on the system. Yonatan Zunger, head of infrastructure at Google, has called this potential threat “the most demonically clever computer security attack I’ve seen in years” [24].

B. Detecting counterfeit devices

Since various motivations drive IC counterfeiting, the signals of counterfeit microelectronics also vary. For example, a simple substitution of one part for another is easily detectable by verifying that the part number printed on the IC is as expected. Mis-labeled parts may also be easily verified by electrical tests, especially if the two parts have completely different functionality (as opposed to, for example, having the same functionality but different tolerances). Recycled parts may show visible external signs of wear on the packaging. Many counterfeits are created by sanding off or painting over the markings of another IC and then adding a new, misleading silk-screened label. These devices can often be detected by looking for signs of uneven finishes, use of abrasive on the IC’s external surface, or inconsistencies or imperfections in the silk-screened markings that may signal that they were not printed by an established manufacturer.

For other types of counterfeiting, more sophisticated examination of an IC is required. For example, the external markings of a counterfeit produced by a determined adversary may be indistinguishable from those of a legitimate part. Particularly worrisome are trojan circuits added to ICs, since they may not affect the IC’s behavior until activated by some very specific sequence of events [43]. For these counterfeits, detailed analysis of the silicon surface – the actual circuits inside the IC – may be required.

III. VISUAL INSPECTION OF MICROELECTRONICS

Among the variety of different types of tests that can be used to verify electronic devices, visual inspection (of either the external packages or internal circuits) has the advantage that it can detect a wide variety of defects and counterfeits. Traditionally, visual inspection is performed by human experts who are trained to identify common visual elements and patterns, and spot important anomalies. Unfortunately, this means that while the microscopy images can be collected efficiently and automatically, actual counterfeit detection generally depends on human inspectors to manually analyze the resulting imagery. This presents obvious scaling challenges, given the huge number of devices in use and their complexity. In this section, we review background on visual inspection of microelectronics, including its efficacy and limitations.

A. Inspection by human experts

If our goal is to avoid labor-intensive manual inspection of large-scale microelectronic imagery, it is important to first

understand the capabilities and limitations of humans on this task. Of course, humans have very powerful cognitive abilities to reason about the visual information they observe in the world. Although computer vision is improving rapidly, and can, on some tasks, outperform humans in limited contexts [27], [29], no algorithms come close to matching human visual inference abilities in general. A practical application of the power of humans to make sense of complex visual information is the use of CAPTCHAs [5] to identify human users online by testing the limits of a remote requester’s ability to recognize text in the face of complicated distractors and transformations.

But while humans have a remarkable ability to process visual information, our reasoning and cognition has important limitations, especially in large-scale inspection contexts. People have a limited processing rate and limited attention span; in particular, humans are known to perform poorly in tasks where they are trying to identify very rare events, such as weapons in baggage or, in our case, anomalies among perhaps thousands of unremarkable examples. Moreover, people make subjective and unconscious inferences about visual information, and can be easily confused by complicated or unexpected visual stimuli. This causes people to perceive scene and object features that do not exist, or to miss important features that do exist, as is demonstrated by well-known optical illusions [11].

Another limitation is humans’ sensing abilities. Of the entire electromagnetic spectrum, humans can only sense light from 380 nanometers (marginally above infrared) to 740 nanometers (well below ultraviolet) [44], so instruments that use these spectra must produce visualizations that are mapped into the visible spectrum. Limits of visual acuity permit humans to only resolve objects around 0.1mm, approximately the diameter of a human hair, from a typical reading distance [44]. But this seemingly tiny distance can hold thousands of 14nm transistors! Humans thus must rely on microscopes and other devices to overcome the physical limitations of their eyes.

B. Visual inspection

Visual inspection is required in a wide variety of fields, from surveillance and security to industrial and quality control applications, and so a variety of papers have studied various inspection tasks and how to make them more successful. For example, Sandia National Laboratories recently published an overview of 212 visual inspection papers from 1958 to 2012 [36]. Most of these papers studied applications where inspectors sought to identify defects, and in particular in situations when defects were rare but the cost of missing them was high.

The survey suggests that there are two broad dimensions along which visual inspection tasks can be characterized. The first axis characterizes the type of information that must be processed by the inspector. Rasmussen published a skilled performance taxonomy in 1983 known as ‘Skills-Rules-Knowledge (SRK),’ which can be used to characterize the degree of difficulty of a visual inspection task [34]. The taxonomy characterizes tasks according to three levels. *Skill-*

level information is perceived as continuous “raw” time-space signals without semantic meaning. Processing this level of information requires the least effort and essentially occurs unconsciously. *Rule-*level information is evidence that helps select or modify rules through past experience, and requires an intermediate degree of effort. *Knowledge-*level information is more semantic and symbolic, and is used for reasoning; these tasks require the highest degree of effort [34]. The other axis describes the level of discrimination required for a given task. Not surprisingly, tasks that require a finer grain of discrimination increase both cognitive load and error rate [18]. Many inspection tasks require challenging, often expert, levels of discrimination.

These two dimensions can be used to assess the anticipated difficulty of a proposed inspection task. For example, suppose a human operator is inspecting circuit board assemblies, using an imaging tool that highlights areas of de-lamination. Since areas of de-lamination are obvious, the operator can identify them with minimal thought. This task requires only a *skills-*level of performance under the SRK taxonomy and does not require a high level of discrimination to identify defects. On the other hand, inspecting memory chips for evidence of tampering requires spotting rework on the leads, evidence of sanding on the package, and subtle modifications of marking. This task is significantly more difficult because it requires *Knowledge-*level performance under the SRK taxonomy. The level of discriminatory skill will likely depend on the type of defect: natural variation that follows statistical distributions might be relatively easy to spot (e.g., accidental device damage), whereas variation caused by an active adversary who is trying to avoid detection likely requires a high level of discrimination skill.

Unfortunately, inspection tasks for microelectronic applications are growing more and more difficult with time along both axes as devices become increasingly complex and sophisticated. Moreover, and counter-intuitively, humans’ ability to spot defects and other anomalies becomes worse as the number of anomalies decreases: Harris found that inspection accuracy suffers as the defect rate decreases [26], while Megaw observed that regardless of the actual number of defects in an observation sample, inspectors appear to search for approximately five [32]. Visual inspection performance varies based on a person’s age, visual acuity, intelligence, aptitude, personality, experience, visual lobe, scanning strategy, and biases [36]. If the frequency of defects changes over time, biases will be induced in the inspectors [36]. Drury found that error rates of 20% to 30% are common across multiple types of inspection tasks [19]. Finally, inspectors can only focus on inspection tasks for a small period of time before they fatigue. Teichner found that vigilance significantly degrades just 15 minutes into a visual inspection task [39], while Drury and Fox find that defect detection could deteriorate 40% in 30 minutes [19]. Finally, from an ethical and quality-of-life point of view, inspection tasks are inherently stressful for inspectors [36].

Given these limitations, much work has studied how to

maximize the efficacy, accuracy, and well-being of inspectors. For example, Drury and Watson demonstrated that inspectors are better at classifying (i.e., deciding if a suspect feature is defective or not) than searching (i.e., scanning for and identifying suspect regions) [21]. The limitations of a single human inspector can be reduced by having multiple inspectors work together. Drury et al. explored five methods to incorporate two inspectors to increase accuracy: (1) each inspector inspects half of the batch of items in parallel, (2) both people inspect every item and accept an item only if both inspectors accept it, (3) both people inspect every item and reject it only if both reject it, (4) the second inspector only inspects items accepted by the first, and (5) a second inspector inspects only items rejected by the first. They concluded that errors are minimized when both inspectors inspect every item and reject an item only if both reject it [20].

C. Imaging instruments

Some of the above limitations of the human visual system can be corrected by imaging instruments. Microscopes allow us to see individual red blood cells that are 10 micrometers wide, bacteria that are 1 micrometer wide, and even individual viruses that are just 30-300 nanometers wide [9]. However, microelectronic device components are even smaller. To manufacture, troubleshoot, and evaluate microelectronics, the semiconductor industry relies on a variety of microscopes that fall under three general categories: light microscopes, electron microscopes, and scanning probe microscopes. This section will briefly survey these categories of instruments with a focus on tools that are commonly used for inspecting semiconductors.

The first microscopes used visible light and quickly advanced to the limits of diffraction. Since light microscopes image using the visible spectrum, they have a useful magnification of up to about 1000 times, which permits resolving visual features of about 100 nanometers [22]. Traditional compound light microscopes are relatively inexpensive, but since IC components are typically opaque to the visible wavelengths, using traditional light microscopes to inspect microelectronics typically involves destroying samples (e.g., physically removing them from the packaging). In contrast, Near-field Scanning Optical Microscopes use a laser light source that illuminates the component under inspection while a sensor probe collects data in a raster scan. These devices offer two distinct advantages for the imaging of microelectronics: (1) they overcome the diffraction limit, allowing increased magnification, and (2) bulk silicon is transparent to the infrared laser light source, allowing imaging of active electronics through the backside of the die.

Electrons have a wavelength approximately 1000 times shorter than visible light, allowing for greatly enhanced magnification. Today's most powerful microscope is the Transmission Electron Microscope, which has a magnification factor of approximately 1,000,000, offering a resolving power on the order of 1 Angstrom [22]. These microscopes can image through very thin cross sections of material and thus are generally

limited to destructive inspection of microelectronics, requiring a potentially time-consuming preparation (e.g. slicing) of the sample ahead of time.

A Scanning Electron Microscope functions in a similar way to a Near-field Scanning Optical Microscope, and is particularly useful for high-magnification imaging of surfaces. The maximum magnification of modern instruments is approximately an order of magnitude less than Transmission Electron Microscopes, but still generally requires a pre-imaging preparation process that destroys the sample under test.

Several other types of devices may be useful for microelectronic inspection. For example, Photon Emission Microscopes measure the emission of light (photons) that are radiated in active electronics [22]. These instruments operate much like Scanning Optical Microscopes described above, and require topside access to the die, which requires a destructive sample preparation process. In contrast, X-ray microscopes image with short-wavelength X-rays, providing illumination that penetrates semiconductor packaging and thus is non-destructive. They provide a resolution of about 10 nanometers. Another device useful for non-destructive inspection is the acoustic microscope, which operates much like an X-ray microscope but uses long-wavelength ultrasonic acoustic energy. They are most useful for low-magnification inspection of IC packaging. Other scanning probe microscopes include magnetic, capacitive, and atomic force variants that operate similarly, while measuring different physical phenomena. Each of these offers a unique, and very specific, capability that is limited to destructive imaging applications.

D. Computational microscopy

Despite the limitations discussed in the last section, visual inspection remains in widespread use for inspection and analysis of microelectronics, including for identifying counterfeit parts. To alleviate these limitations, an accepted best practice is that when evidence of counterfeit parts is discovered, the entire lot of parts is rejected. This reflects the assumption that when counterfeit parts are introduced into the supply chain, the part lots are often homogeneous (i.e., all parts are bad). Thus the discovery of only one suspect part is needed, which significantly relaxes the pressure on any single inspector to make the correct judgment about any single part.

However, these assumptions may not hold for critical applications, where a motivated malicious adversary may be actively trying to prevent counterfeit identification. Heterogeneous part lots that contain both legitimate parts and those with added malicious functionality present a much more difficult problem to human inspectors, for example. Finding the subtle changes in a circuit that indicate a malicious modification typically requires fine-grained analysis by microscope, but microscopes create large amounts of image data, which increases the burden on the human inspector.

One possible approach to alleviate this problem is to use automated image analysis. Of course, computer vision techniques have long been used for visual inspection tasks [10]. These applications tend to be most successful when the environment is

highly constrained and the set of possible scenarios is limited, such as in assembly line contexts, so that they require just Skills-level reasoning under the SRK taxonomy introduced in Section III. Machine vision-based inspection typically consists of comparing a test image against ground truth, and measuring the difference between the images, where a “passable” image is defined to be visually similar to the ground truth up to some quantified threshold.

In these scenarios, computer vision can match or outperform human-level inspection abilities. For example, Carrasco et al. described the use of computer vision to look for defects in glass bottles where the accuracy exceeds that of human operators [12]. Other work, largely in biology and other physical sciences, has developed successful techniques for computer vision on microscopy images [23], [46].

In contrast, there is relatively little prior published work that applies computer vision to microelectronic image analysis. Mahmood et al. detected counterfeit parts in real time using X-ray microscopy [30]. They used local binary pattern (LBP) features with both support vector machines (SVMs) and deep learning classifiers. Subsequent work made improvements by lowering the X-ray dose during image acquisition [31]. Zhou et al. detected hardware trojans inserted during fabrication by embedding visible watermarks in fill cells [45], although this approach requires access to and control over the design prior to fabrication. Adato et al. addressed this limitation by proposing laser confocal microscopy to map ICs by classifying logic gates [8]. Cilingiroglu et al. proposed overcomplete dictionaries to evaluate microelectronic imagery [13]–[15], inspired by the bag-of-visual-words approach that is popular in other computer vision problems [17], [42].

One reason for the limited work on automatic microelectronics microscopy analysis is that fine-grained and less well-defined tasks, such as identifying a counterfeit of an arbitrary electronic device, require greater sophistication; under the SRK taxonomy, these more challenging tasks require Rules or Knowledge-level abilities. For example, real-world imaging of arbitrary devices introduces visual clutter, illumination difficulties, noise, and occlusion that must be ignored or accounted for during the inspection process. Of course, despite decades of very active research and many significant breakthroughs, computer vision is nowhere near human capability to perform general visual tasks.

However, the last several years have brought exciting advances in computer vision, driven largely by deep machine learning [28]. The basic idea behind these approaches is that, instead of using manually-designed algorithms to extract features from images and then feeding them into a separate classifier (such as an SVM), the learning algorithm takes raw image pixels as input, and learns the ideal feature representations along with the high-level classifier in a single unified training procedure. When successful, this approach produces significantly better results than the traditional manually-designed features. However, it also introduces several challenges. First, large-scale training data is typically required, on the order of tens of thousands to millions of images [28]. While such

data is readily accessible for some applications such as web images, it is difficult to imagine collecting a dataset of millions of counterfeit ICs, for example. Due to this need for large-scale training data, the learning algorithms also require large amounts of computation, typically in the form of high-end graphics processing units (GPUs). But perhaps most problematic is that deep learning-based techniques are largely “black boxes,” making decisions without providing much insight into why or how the decision was made (although much current research is focusing on addressing this limitation [37]).

IV. TOWARDS AUTOMATED INSPECTION OF INTEGRATED CIRCUITS

Despite the limitations of computer vision and deep machine learning discussed in the last section, there have been so many recent successes on such a wide variety of applications and problems that it is natural to investigate deep computer vision-based techniques for microelectronic image analysis. Although our eventual goal is to use computer vision for analyzing integrated circuits at the microscopic scale, here we start with an initial step of matching ICs based on visual features of their external packages. Analyzing external features could be useful in and of itself, however, since a large fraction of counterfeits are apparent based on packaging anomalies alone [38].

A. Dataset

We are not aware of any publicly-available, large-scale datasets of counterfeit microelectronics and do not have ready access to such devices. As a proxy, we instead collected a large-scale and realistic collection of images of IC packages, and tested computer vision-based methods for finding similarities and correspondences across this set. We imagine a scenario in which a large number and variety of microelectronic devices are entering some part of the supply chain, and we want to organize this unstructured collection by identifying groups of similar parts, and quantifying the visual similarities and differences between them.

We collected our dataset by downloading a seed set of about 150 printed circuit board (PCB) images from [33]. We then used this seed set to query for more images using Google Image Search, which tries to find web images that are visually similar to a given target image. This generated a collection of about 1,500 images. We manually removed images that were of low resolution, that were not real photos of actual printed circuit boards (e.g., cartoons and drawings), that contained prominent watermarks, that were duplicates, etc., to produce a final set of 500 images.

We then manually marked bounding boxes around each IC in each of the images, excluding discrete electrical components (e.g., discrete transistors, capacitors, transformers, etc.), as well as empty IC sockets. The bounding boxes were made large enough to include the full IC packages including any visible electrical pins or leads. We cropped out each IC bounding box and then manually rotated it, if necessary, so that any text on the IC was right-side-up. We used the open-

source, web-based LabelMe tool [35] to collect the bounding box annotations.

We then used Amazon Mechanical Turk [6] to transcribe the textual markings on a subset of 814 of the ICs. In particular, we presented each IC to two Mechanical Turk users, each of whom was asked to indicate whether any text was visible on the IC, and if so, to fill in three text fields whenever possible: (1) the IC’s country of origin, (2) the IC’s manufacturer, and (3) any ID numbers or other textual labels visible on the IC’s surface. When text markings were only partially visible, we asked users to type in as much as they could see and to enter illegible characters as asterisks. To help ensure high-quality labels, we required Mechanical Turk users to have “Master” status, meaning that they have a long track record of work on the site. After the labeling process was complete, we compared the two independent transcriptions produced for each IC, and manually resolved any disagreements between them. In total, 34 Mechanical Turk users participated in the labeling process.

Figure 1 shows a small sample of the dataset, including a PCB image with IC ground truth locations superimposed as bounding boxes, and the resulting cropped images of individual ICs.

B. Matching experiments

Our goal in these preliminary experiments was to apply a computer vision algorithm to automatically identify similarities among the 818 integrated circuit images. This task is a first step towards identifying an IC based on its external packaging, which in turn could be used for finding mislabeled or counterfeit parts in a large-scale inspection or manufacturing environment. The task of comparing two images may seem easy from the perspective of a human observer, but differences in image alignment, illumination, perspective, etc. make it a non-trivial task for an automatic algorithm. On the other hand, finding similarities amongst nearly a thousand images would be time consuming if done by hand. Thus while this task is designed to be simple and preliminary, it already suggests the value of automated techniques for comparing and organizing large-scale IC image collections.

To perform the image clustering, we applied a recent technique in computer vision that was designed to find optical flow (motion patterns) between two adjacent frames of a video sequence [41], adapting it to instead measure the visual similarity between two static images. In particular, we first convert each image to grayscale and resample it to 250×250 pixels, regardless of original aspect ratio. We then run the deep matching algorithm of Weinzaepfel et al. [41], which estimates the optical flow between the two images; in other words, for each small patch in one image, it finds the most similar patch in the second image. We take these correspondences and compute the amount of error between the two images, yielding a measure of their similarity while still allowing for some flexibility to account for changes in illumination, perspective, noise, etc.

Figure 2(a) presents a visualization of a subset of 100 IC images. The table compares each of the 100 images to each

of the other 100 images, where a blue cell indicates high similarity and red indicates low similarity. Of course, the greatest similarities lie along the diagonal, corresponding to images being compared to themselves. There are three other prominent clusters of blue (located near the diagonal in the middle and bottom-right portion of the table). The images in each of these clusters are shown in Figure 2(b). As expected, we see that the clusters correctly correspond to similar or identical ICs.

While these results are just a first step, they show that this technique could be used to quickly find groups of similar and nearly-similar ICs in a large unorganized collection, a task that would be very labor-intensive for a human. Importantly, the algorithm has no prior knowledge about the appearance of these ICs and does not use text recognition or other means; it simply attempts to find visual correspondences in an unsupervised way. This is in contrast to existing techniques for IC analysis, which instead compare images to known reference (“gold standard”) exemplars. For example, our technique could be used in scenarios where we may need to identify and organize large lots of ICs but we do not know ahead of time which ICs will be in the group.

V. CONCLUSION AND DIRECTIONS FOR FUTURE WORK

In this paper we highlighted the problems posed by supply chain threats directed against integrated circuits. One possible technique for mitigating this threat is through visual analysis of microelectronic devices, including inspecting both external packages and internal components at a microscopic scale. Given their unprecedented size and complexity, however, conducting this visual analysis by hand is infeasible. Automating this analysis process through computer vision is a promising solution, especially given the impressive performance it has achieved in other applications. We surveyed the general problem area, laid out a research agenda for the computational microscopy of microelectronics, and presented an initial dataset and preliminary IC matching results.

More broadly, we hope future work will study applying computer vision to microelectronic imaging for a variety of tasks including (1) detecting defects and other issues that impact part quality or reliability, (2) identifying suspect counterfeit parts, and (3) making inferences about the functionality of ICs. This problem space is enormous given the large variety of microelectronic part types, the numerous imaging modalities (e.g., optical, electron), the level of expertise required for fine-grained inferences about IC images, and the mind-boggling complexity and scale of modern semiconductors.

ACKNOWLEDGMENTS

We thank Katherine Spoon for assisting with the Mechanical Turk study. At Indiana University, this paper is part of a broader project including Sven Bambach, Jacob Beauchamp, Saúl Blanco, Joshua Cannon, Ben Lewis, and Jacob Nixon, and we thank them for helpful discussions and other assistance. This research was sponsored by the Naval Engineering Education Consortium (NAVSEA) contract N00174-16-




Fig. 1. A sample from our dataset, showing (a) one of the printed circuit boards with ground truth bounding boxes around integrated circuits superimposed, and (b) the individual cropped ICs. The dataset also includes ground truth transcriptions of text appearing on the ICs (not shown).

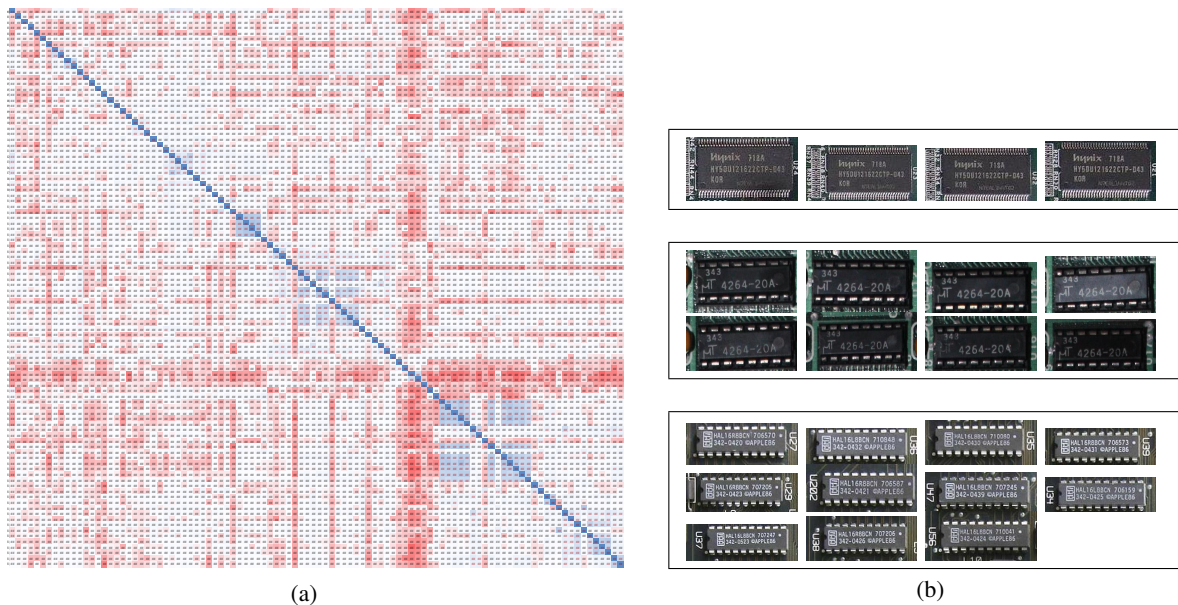


Fig. 2. Preliminary IC matching results. (a) Visualization of similarity scores for a subset of our dataset (100 IC images), where blue indicates greater similarity and red indicates lower similarity. The strong blue line along the diagonal corresponds to high self-similarity scores. (b) The IC images in the three prominent blue clusters of (1).

C-0016, with support of NSWC Crane Division in Crane, Indiana. It used compute facilities donated by Nvidia, Inc., and the Romeo FutureSystems Deep Learning facility, which is supported in part by Indiana University and the National Science Foundation (RaPyDLI-1439007).

REFERENCES

- [1] <http://www.semiconductors.org/>.
- [2] <http://www.gpo.gov/fdsys/pkg/BILLS-112hr1540enr/pdf/BILLS-112hr1540enr.pdf>.
- [3] <http://www.ti.com/corp/docs/kilbyctr/jackbuilt.shtml>.
- [4] <http://www.purdue.edu/newsroom/research/2012/120219KlimeckAtom.html>.
- [5] <https://www.google.com/recaptcha/intro/index.html>.
- [6] <http://www.mturk.com/>.
- [7] Winning the battle against counterfeit semiconductor products. Technical report, Semiconductor Industry Association, 2012.
- [8] R. Adato, A. Uyar, M. Zangeneh, B. Zhou, A. Joshi, B. Goldberg, and M. S. Unlu. Rapid mapping of digital integrated circuit logic gates via multi-spectral backside imaging. *arXiv:1605.09306*, 2016.
- [9] T. Allen. *Microscopy: A Very Short Introduction*. Oxford University Press, 2015.
- [10] J. Beyerer, F. P. Leo, and C. Frese. *Machine Vision: Automated Visual Inspection: Theory, Practice and Applications*. Springer, 2016.
- [11] C.-C. Carbon. Understanding human perception by human-made illusions. *Frontiers in Human Neuroscience*, 8:566, 2014.
- [12] M. Carrasco, L. Pizarro, and D. Mery. Visual inspection of glass bottlenecks by multiple-view analysis. *International Journal of Computer Integrated Manufacturing*, 23(10):925–941, 2010.
- [13] T. Cilingiroglu, A. Tuysuzoglu, W. Karl, J. Konrad, M. Ünlü, and B. Goldberg. Dictionary based image enhancement for integrated circuit imaging. In *International Conference on Acoustics, Speech and Signal Processing*, 2013.
- [14] T. Cilingiroglu, A. Uyar, A. Tuysuzoglu, W. Karl, J. Konrad, B. Goldberg, and M. Ünlü. Dictionary-based image reconstruction for super-resolution in integrated circuit imaging. *Optics express*, 23(11):15072–15087, 2015.
- [15] T. Cilingiroglu, M. Zangeneh, A. Uyar, W. Karl, J. Konrad, A. Joshi,

- B. Goldberg, and M. Unlu. Dictionary-based sparse representation for resolution improvement in laser voltage imaging of cmos integrated circuits. In *Design, Automation & Test in Europe Conference*, 2015.
- [16] M. Crawford, T. Telesco, C. Nelson, J. Bolton, K. Bagin, and B. Botwin. Defense industrial base assessment: Counterfeit electronics. Technical report, U.S. Department of Commerce, 2010.
- [17] G. Csurka, C. Dance, L. Fan, J. Willamowski, and C. Bray. Visual categorization with bags of keypoints. In *ECCV Workshop on Statistical Learning in Computer Vision*, 2004.
- [18] C. Drury and J. Addison. An industrial study of the effects of feedback and fault density on inspection performance. *Ergonomics*, 16(2):159–169, 1973.
- [19] C. Drury and J. Fox. The imperfect inspector. *Human reliability in quality control*, pages 11–16, 1975.
- [20] C. Drury, M. Karwan, and D. Vanderwarker. The two-inspector problem. *IIE Transactions*, 18(2):174–181, 1986.
- [21] C. G. Drury and J. Watson. Human factors good practices in borescope inspection. Retrieved July, 8:2002, 2001.
- [22] R. F. Egerton. *Physical Principles of Electron Microscopy: An Introduction to TEM, SEM, and AEM*. Springer, 2016.
- [23] D. Fernandez, R. Bhargava, S. Hewitt, and I. Levin. Infrared spectroscopic imaging for histopathologic recognition. *Nature biotechnology*, 23(4):469–474, 2005.
- [24] S. Greengard. Are computer chips the new security threat? *Communications of the ACM*, 60(2):18–19, 2017.
- [25] S. Greengard. The future of semiconductors. *Communications of the ACM*, 60(3):18–20, March 2017.
- [26] D. Harris. Effect of defect rate on inspection accuracy. *Journal of Applied Psychology*, 52(5):377, 1968.
- [27] K. He, X. Zhang, S. Ren, and J. Sun. Delving deep into rectifiers: Surpassing human-level performance on ImageNet classification. *arXiv 1502.01852*, 2015.
- [28] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *Neural Information Processing Systems*, 2012.
- [29] C. Lu and X. Tang. Surpassing human-level face verification performance on LFW with GaussianFace. In *AAAI Conference on Artificial Intelligence*, 2015.
- [30] K. Mahmood, P. L. Carmona, S. Shahbazmohamadi, F. Pla, and B. Javidi. Real-time automated counterfeit integrated circuit detection using x-ray microscopy. *Applied Optics*, 54(13):D25–D32, 2015.
- [31] A. Markman and B. Javidi. Integrated circuit authentication using photon-limited x-ray microscopy. *Optics Letters*, 41(14):3297–3300, 2016.
- [32] E. Megaw. Factors affecting visual inspection accuracy. *Applied ergonomics*, 10(1):27–32, 1979.
- [33] C. Pramerdorfer and M. Kampel. A dataset for computer-vision-based pcb analysis. In *Machine Vision Applications*, 2015.
- [34] J. Rasmussen. Skills, rules, and knowledge; signals, signs, and symbols, and other distinctions in human performance models. *IEEE transactions on systems, man, and cybernetics*, 1(3):257–266, 1983.
- [35] B. C. Russell, A. Torralba, K. P. Murphy, and W. T. Freeman. LabelMe: a database and web-based tool for image annotation. *International Journal of Computer Vision*, 77:157–173, 2008.
- [36] J. E. See. Visual inspection: A review of the literature. Technical Report SAND2012-8590, Sandia National Laboratories, 2012.
- [37] R. R. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. *arXiv:1610.02391*, 2017.
- [38] M. M. Tehraniipoor, U. Guin, and D. Forte. *Counterfeit integrated circuits*. Springer, 2015.
- [39] W. H. Teichner and M. J. Krebs. Laws of visual choice reaction time. *Psychological review*, 81(1):75, 1974.
- [40] J. Villasenor. Compromised by design? Securing the defense electronics supply chain. Technical report, Brookings Institute, November 2013.
- [41] P. Weinzaepfel, J. Revau, Z. Harchaou, and C. Schmid. Deepflow: Large displacement optical flow with deep matching. In *International Conference on Computer Vision*, 2013.
- [42] J. Xiao, J. Hays, K. A. Ehinger, A. Oliva, and A. Torralba. Sun database: Large-scale scene recognition from abbey to zoo. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2010.
- [43] K. Yang, M. Hicks, Q. Dong, T. Austin, and D. Sylvester. A2: Analog malicious hardware. In *IEEE International Symposium on Security and Privacy*, 2016.
- [44] M. Yanoff and J. Duker. *Ophthalmology*. MOSBY Elsevier, 2009.
- [45] B. Zhou, R. Adato, M. Zangeneh, T. Yang, A. Uyar, B. Goldberg, S. Unlu, and A. Joshi. Detecting hardware trojans using backside optical imaging of embedded watermarks. In *Proc. Annual Design Automation Conference*, 2015.
- [46] S. Zhou. *Medical Image Recognition, Segmentation and Parsing*. Academic Press, 2015.