Physical Object Protection based on Micro-structure Fingerprinting

Dr Fokko Beekhof

Université de Genève

Physical Object Protection based on Micro-structure Fingerprinting

Université de Genève

F. P. Beekhof

1 Introduction

Users and manufacturers of physical objects face problems due to uncertainty about the origin of goods, which are physical objects. Counterfeit products may be confused with authentic items, incurring a loss to manufacturers in terms of money and reputation of the brand. Consumers may receive products of lower quality and value than what they paid for, and in certain cases, such as counterfeit pharmaceutical products, there are potential health risks. Counterfeiting manifests itself in several forms, such as the production of imitations, rebranding recycled items as new products, illegal franchising, or any combination thereof. The problem of counterfeit electronics is even widespread in military products [8, 13]. At the University of Geneva, we have investigated the problem of protecting goods together with our industrial partners Anteleon Imaging and u-nica systems.

Historically, ways to secure products include the use of special inks [2], anti-copying visible patterns [14], embedded holograms or microtext [17, 19]. However, the protection mechanism behind these techniques requires that the object under protection must be modified in some sense. Proprietary protection in physical forms, such as a hologram or RFID chip must be embedded into, or attached onto the object, which may not be physically possible, or simply undesirable from a commercial or legal standpoint. If each object requires a uniquely distinguishable identity, then the hologram, microtext or chip attached to each object must be unique as well, which can lead to increased cost.

1.1 Micro-structures

An alternative solution is to use unique features that can not be recreated by a forger, or only at prohibitive cost, and are found on almost all objects. Microstructures meet all these conditions. The surfaces of most of the objects surrounding us naturally exhibit a random structure at a small scale, as shown in Figure 1, and are visible by looking closely at an object. These random structures are known as microstructures. The random pattern of microstructures is largely determined by the physical properties of the material, be it paper, metal, plastic, etc, and the microstructures have a random character and are unique for almost all items. These properties make microstructures an ideal candidate for recognizing individual objects, much like the human fingerprint or iris can be used to identify people. By registering all original items in a database during the production stage, it becomes possible to determine if an object of unknown origin is authentic.

Although microstructure images contain all the information necessary to recognize objects, it is cumbersome to work with whole microstructure images, certainly if thousands of objects are registered. Instead, we have turned to digital content fingerprinting, a technique that transforms digital content, such as images, audio or video, into a short sequence of bits. These sequences of bits are small tokens that are used to identify different contents, which are then referred to as fingerprints. They take their name from human fingerprints, which fulfil a similar role. There is a large body of literature available on digital content fingerprinting.

An important contribution has been made by Fridrich and Goljan [5, 6, 7], who proposed a robust hashing function for images based on projections onto blurred randomly generated hyperplanes. Inspired by this work, the use of such random projections is the subject of a large part of the research in the Stochastic Information Processing group, and forms the basis of the fingerprint scheme used in the context of my Doctoral Thesis.

1.2 Requirements for fingerprinting

In most earlier work on digital content fingerprinting, the main concern was to develop an algorithm and demonstrate empirically that common signal processing operations and addition of noise to the image did not lead to a large number of altered bits in the fingerprint. The requirements for fingerprint techniques stated in different publications from this era tend to overlap, but without a coherent view. There is a vast body of literature that does not explicitly mention any design goals of fingerprinting, but simply reports empirical results for tests of robustness against a variety of common attacks such as cropping, JPEG compression, scaling etc. [10, 11, 15, 18]. The results are measured by counting how many bits in a fingerprint change when the image is modified.

Doets [3], who analyzed the audio fingerprinting scheme of Kalker and Haitsma of Philips Electronics[9], was the first to propose a systematic and comprehensive list of requirements for fingerprinting. However, the formulation of the requirements could still be improved, and to do so, I turned to Information Theory.

Information Theory is a branch of mathematics that came into existence in 1948, in a now-famous publication of Claude Shannon at Bell Labs about communication [16]. Although Shannon's work on Information Theory stretched far beyond communication and includes compression and cryptographic principles, the work on communication is most significant in our research. It allows one, for example, to express limits on the number of items that can be protected for a given amount of noise in the images. The most important basic concepts in information theory are the entropy of a random variable, which, for simplicity can be seen amount of information that the variable contains; and the mutual information between two variables, which, loosely speaking, can be seen as the amount of information that is "shared" between the two.

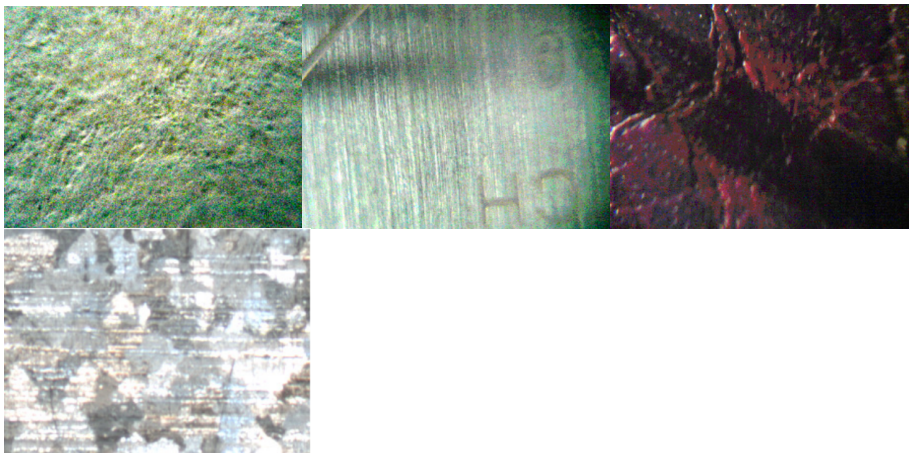

(a) Paper (b) Metal (c) Leather (d) Aluminium

Figure 1: Examples of micro-structures

The requirements formulated in my work are largely compatible with Doets' requirements, but using Information-Theory allows a more correct and stricter evaluation of fingerprinting schemes, and shows the relations between the image and fingerprint, as well as between the requirements in a provable and quantifiable way. This means that the behaviour of a fingerprinting scheme can now be proven, rather than only empirically tested, although empirical tests are still useful to confirm mathematical predictions. The most important requirements are explained in the following paragraphs.

Informativeness refers to the amount of information of the microstructure image that is retained in the fingerprint. Existing work frequently demands that fingerprints must contain a lot of information, but such a formulation does not ensure a link between the information in the fingerprint and the microstructure. It has so far been overlooked that a high information-content or entropy can simply be due to certain randomizing influences, for example by using a secret key. The new formulation stipulates that only the information in the fingerprint that is derived from the microstructure counts,

and is based on the mutual information between the microstructure image and the fingerprint.

Compactness is a measurement of efficiency of storage of the fingerprints. Information-theory is only concerned with the amount of actual information, but does not concern itself with how dense this information is written down. Much like the same text can require more or less characters depending on which language is used to write it down, the fingerprints can require different amounts of storage space, depending on their representation - even when they contain, in an information-theoretic sense, the same amount of information. The compactness of the representation can be measured relative to the minimal amount of storage space required for all information in the fingerprint. If, for example, a fingerprint contains 20 bits worth of information, the storage space should be as close as possible to 20 bits.

Robustness indicates how much of the information about the microstructure is retained in the fingerprint if the microstructure image is noisy. Ideally, the fingerprint should retain as much information as is available, or at least lose information at most proportionally to the loss of information between the microstructure images for registration and verification. The robustness is expressed as the mutual information of a fingerprint taken from a near- perfect microstructure image, assumed to be taken during the registration, and fingerprint derived from a noisy microstructure image, taken for verification purposes.

Distinctiveness describes the difference between fingerprints taken from images of different microstructures. Although the informativeness and robustness ensure that there is information about the microstructure image in the fingerprint, even when the the image is noisy, this is not sufficient to tell objects apart based on their fingerprints. The distinctiveness of fingerprints requires exactly this: we must be able to recognize different object by their fingerprints, so the fingerprints must be sufficiently different.

Resilience expresses the ability of the system to maintain acceptable performance, in terms of the criteria mentioned before, when a malicious party actively attempts to subvert the system. It differs from the previously stated requirements because it describes the behaviour of other requirements, and can therefore be seen as a “meta-requirement”. The concept of active attacks on forensic systems, and thus resilience is a new and exciting direction of research. Universality is another “meta-requirement”, by which we mean to say that the stated requirements should be met regardless of certain statistics, particularly those of the noise. Universality is a known property that is desirable in many domains, such as compression [23, 24].

The importance of the reformulation can be seen from the example of privacy amplification, which is a technique where some of the fingerprint’s ones and zeros are flipped at random, thus destroying a part of the information. The goal of privacy amplification is to reveal less information about the original sample, and has been proposed for biometrics. Biometrics is on a mathematical level a closely related field, and one where the protection of privacy is of critical importance. Following the requirements found in older literature on fingerprints for images, privacy amplification would appear to improve the properties of a fingerprinting scheme: the randomness of the bits will increase and any similarities between fingerprints of different objects will decrease. However, the proposed information-theoretic definitions show that exactly the contrary is the case: the informativeness in actually decreased, and consequently, so is the distinctiveness. Based on the new formulation of the requirements, it becomes clear that it would make more sense to use smaller fingerprints instead of randomizing part of the data. Smaller fingerprints also contain less information about the original content, but require less storage space and processing power.

2 Fingerprint Recognition

Once a fingerprint has been calculated from an object of unknown origin, we can verify its authenticity by comparing it to the list of fingerprints of authentic objects. If a match is found, we assume that the object is authentic. Unfortunately, the noise in the microstructure images will lead to small errors, so some bits will be different in the stored and calculated fingerprint, even if we test an authentic item.

Because of the errors, we can not find the query in the list with a direct lookup system, much like it is very difficult to efficiently find a name with spelling errors in the phone book – imagine what happens if the first letter changes. The naive approach would be to compare the calculated fingerprint with the entire list of stored fingerprints until a match is found, which can become prohibitively slow when the list contains (very) many items. A brand protection system must be able to provide an answer within a few seconds in order to be practically usable.

We propose the following strategy as an alternative to those based on indexing methods [1]. Some information is available about the likelihood of error for each bit. We can use this information to make an informed guess about the bits that are in error, and start changing those bits one by one, creating a tree of variations. Bits that are changed are referred to a branch variables. In Figure 2, the fingerprint from an object under investigation is represented by $\mathbf{b_y}$, and the bits to change, or branch variables can be either flipped from one to zero or vice versa, or kept in their original state. Each black dot represents a new variation of the fingerprint. Although finding a closest match in a database of fingerprints is hard, checking whether or not a single variation exists in the database is very easy and fast.

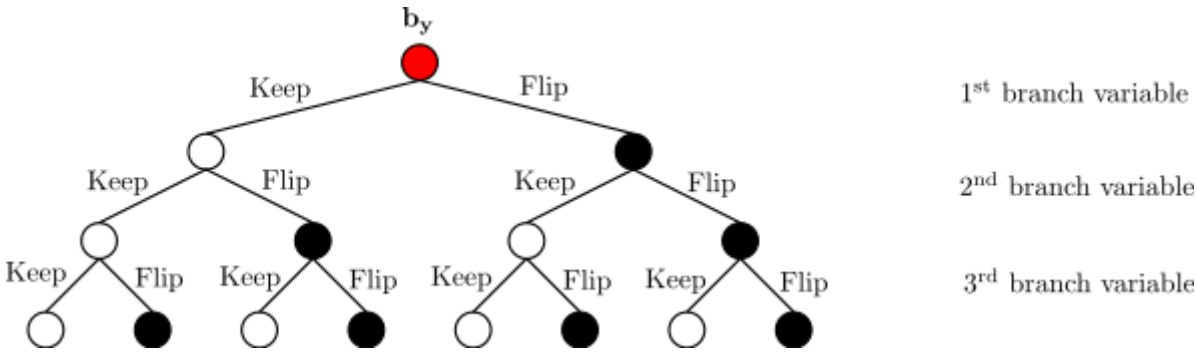

Figure 2: Spanning the search space by flipping bits.

Generating and testing the presence of variations results in a search pattern that eventually includes a perfect match with a fingerprint in the list if the object is authentic. The search pattern through the space of fingerprints is shown in Figure 3, where each black node represents a different fingerprint variation. In the Figure, the fingerprint of an authentic object under investigation is labelled $\mathbf{b_y}$, and the fingerprint of a registered object m is represented by $\mathbf{b_x}$(m) as an example. Because there is a match within the search space, we conclude that the object under investigation must be the authentic object m. The expected distance between the observed and registered fingerprints of an authentic object is $\mathbf{Lp_b}$, where **L** is the length or number of bits in a fingerprint, and $\mathbf{p_b}$ is the probability of a single bit being in error, which can be calculated from the amount of noise in the microstructure images. The search space must be limited to a carefully chosen distance from $\mathbf{b_y}$, the starting point of the search. If no match with the database is found within a chosen distance from the starting point, then we must assume that the object does not match any authentic object, and thus conclude that the object under investigation is counterfeit. Therefore, the limit on the distance from the starting point of the search plays an important role. If it is too short, the system will not find matches for authentic items, but if it is too long, the system might match the fingerprint of a counterfeit item with that of an authentic item. The limit that leads to the least amount of errors can be found mathematically, and depends strongly on the amount of noise in the images. Intuitively, this makes sense: if the microstructures are clear in the images, the system will do well, and if the system must work with low-quality images, it will do less well.

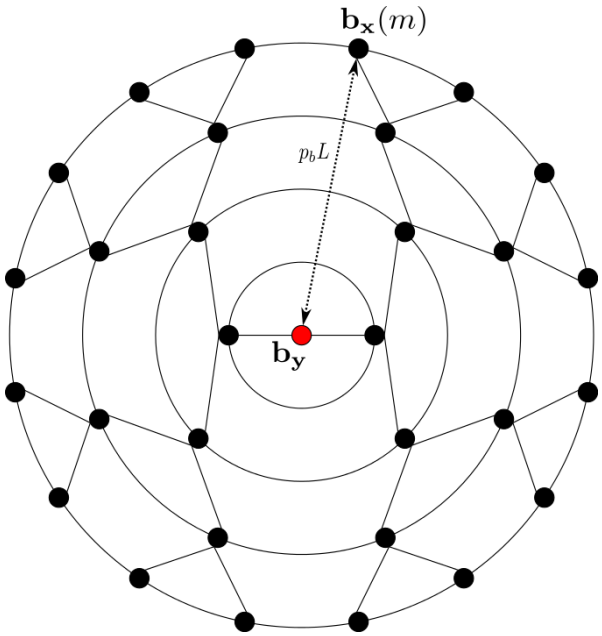

Figure 3: Spanning the search space by flipping bits.

This search strategy is particularly well-suited for large databases and a relatively low noise-level, whereas systems based on indexing are more suited when the microstructure images are of lower quality. Based on an analysis of both approaches, I was able to find the break-even point, which allows system designers to choose the right approach depending on the noise level.

This basic fingerprint search algorithm can be extended using the informa- tion about the reliability of the individual bits. During the calculation of the fingerprint, there is information available about the probability that a particular bit will have changed value due to the noise.

This information can be used to guide the search through the space of fingerprints so that it becomes more efficient, but also to come to an alternative measure that replaces the distance between fingerprints. This “intelligent” measure allows the protection system to decide if an object is authentic or counterfeit with greater certainty.

3 Attacks on Protection Systems

To evaluate the performance of protection systems, one must consider both the probability that the system will mark an authentic object as counterfeit, and that a counterfeit object is declared authentic. The latter is the goal of counterfeiters: that counterfeit objects can pass for originals.

In existing research, it is usually assumed that the fingerprint of a counterfeit object is generated at random, completely independent from the fingerprints of any original contents [4, 12, 20, 21, 22], which is a rather optimistic assumption. I then started to wonder if counterfeiters would be able to defeat the system by creating special structures, even if those were not the original microstructures. This can be expressed as the assumption that counterfeiters are somehow able to produce a counterfeit object which has a fingerprint that is somewhat similar to the fingerprint of one of the authentic objects. It is because of this recent direction of research that the formulated requirements include resilience, the ability to withstand attacks.

The effects of an attack can be shown mathematically. We assume that all bits in a fingerprint behave completely independently from each other. If a fingerprint calculated during the registration phase from a near-perfect microstructure image is compared to the fingerprint of a noisy image of the same microstructure, we expect that a small number of bits, for example 20%, will be different between the

two fingerprints. The number of bit-errors will vary slightly, and the number of bit-errors can be modeled with a Binomial distribution. When two completely unrelated fingerprints are compared, the chance of two bits being equal is 50%. The classic assumption is that fingerprints of counterfeit objects are indeed unrelated to those of registered objects, so existing analyses tend to be based on the assumption that roughly 50% bits between a counterfeit fingerprint and any authentic fingerprint are different. To be more precise, the number of bits that differ between two fingerprints varies and is governed by a Binomial distribution. In Figure 4, several distributions of the number of bit-errors for fingerprints of 32 bits are shown. On the left, with label “Originals” is the distribution for a 20% chance of bit-error as an example of the behaviour of authentic items. On the right, with label “Random” is the distribution corresponding to 50% error, which represents the classic assumption that counterfeit items are completely random and unrelated to original items. The threshold, shown as a vertical line, limits the size of the search space as described in the previous section. We can now see that for some of the originals, the fingerprints will have a number of bit-errors during verification that falls on the right side of the threshold, and thus outside the search space. This means that in this case, some original items would be marked as counterfeit. On the other hand, practically none of the rightmost curve, corresponding to distances between fingerprints of authentic items and unrelated counterfeit items, falls left of the threshold, which would let the counterfeit item to be declared authentic. The system could be improved by moving the threshold to the right so that fewer authentic items would be rejected.

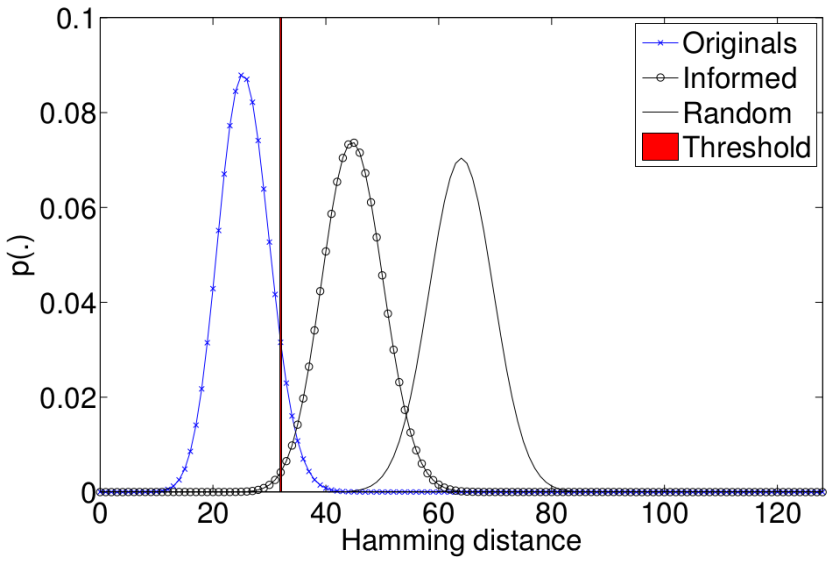

Figure 4: Distributions of number of different bits 20% (authentic), 35% (smart counterfeiting) and 50% (classic assumption of counterfeiting) bit-error probabilities.

The new assumption that counterfeiters can come up with partially matching fingerprints is represented by the curve in the middle, labeled as “Informed”. In this case, a visible part of the curve falls to the left of the threshold, showing that these partial matches have a chance of fooling the system into declaring a counterfeit item authentic.

Whereas the leftmost and rightmost distributions hardly overlap, the leftmost distribution and the one in the middle do. This shows that no matter where the threshold is placed, the system will make mistakes. The consequence of attacks on the brand protection system is thus that the system will be less effective.

3.1 Countermeasures

Fortunately, there is a simple countermeasure that can restore the effectiveness of a brand protection

system: increase the length of the fingerprint. To understand how this works, two ingredients of the previous paragraphs provide the answer. First is the notion that the number of bit-errors can be modeled with a Binomial distribution. Second is the fact that issues arise when the distributions for authentic and counterfeit items overlap. The answer lies thus in an action that changes the distributions so that the overlap between the two is reduced. The Binomial distribution has two parameters, namely the probability $\mathbf{p_b}$ of an event, in this case a bit-error; and **L**, the number of experiments, which, in this case, is equivalent to the number of bits. The probability of a single bit-error $\mathbf{p_b}$ is determined by the noise in the images and the efforts of counterfeiters, and is outside the control of system designers. The length of the fingerprints can be freely chosen by the system designers, and the effects on the overlap for varying fingerprint lengths are shown in Figure 5. Using fingerprints of 32 bits result in a large overlap of the distributions of the numbers of bit-errors. Increasing the fingerprint length to 1024 bits has the effect that the distributions corresponding to authentic and counterfeit items start to separate again, which means that a threshold can be placed such that the system will make very few mistakes.

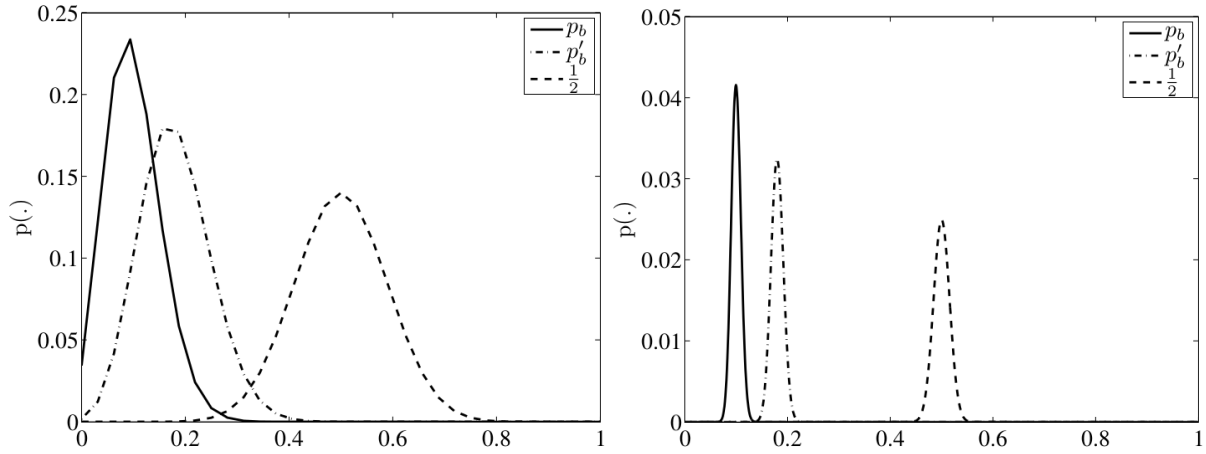

(a) L = 32 (b) L = 1024

Figure 5: Normalized binomial distributions for increasing **L**.

increase of the fingerprint length is likely to increase the amount of information that an attacker can learn about the registered object based on its fingerprint, which can then be used to create more sophisticated attacks. Therefore, the analysis of the trade-off between the information leakage and the separation of the distributions, both as a function of the length of the fingerprint is an interesting open research problem.

4 Conclusions

This article only contains a brief overview of some of the aspects of the research in the domain of brand protection and the underlying problems of fingerprinting techniques, counterfeiting strategies and search algorithms. As the investigation of attacks on brand protection systems by counterfeiters shows, research is progressing and changing to new exciting directions. This kind of research also has value for fields that are related to brand protection, such as biometrics.

5 Acknowledgements

I am grateful to the International Latsis Foundation, the Swiss National Foundation, Anteleon Imaging, and u-nica systems for their support and recognition. Likewise, I would like to acknowledge the help and contributions of my co-workers, notably Prof. Voloshynovskiy, Farzad Farhadzadeh and Maurits Diephuis.

References

[1] Fokko Beekhof, Sviatoslav Voloshynovskiy, Oleksiy Koval, and Taras Holotyak. Fast identification

algorithms for forensic applications. In Proceedings of IEEE International Workshop on Information Forensics and Security, London, UK, December 6–9 2009.

[2] Ingemar Cox, Matthew L. Miller, and Jeffrey A. Bloom. Digital watermarking. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2002.

[3] P.J.O. Doets. Modeling Audio Fingerprints: Structure, Distortion, Capacity.

PhD thesis, Technical University of Delft, 2010.

[4] Farzad Farhadzadeh, Sviatoslav Voloshynovskiy, and Oleksiy Koval. Performance analysis of identification system based on order statistics list decoder. In IEEE International Symposium on Information Theory, Austin, TX, June, 13-18 2010.

[5] J. Fridrich. Robust bit extraction from images. In The International Conference on Multimedia Computing and Systems, volume 2, pages 536– 540, Florence, Italy, June 1999.

[6] J. Fridrich. Visual hash for oblivious watermarking. In P. W. Wong and E. J. Delp, editors, Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, volume 3971 of Society of Photo-Optical Instrumentation Engineers (SPIE) Conference Series, pages 286–294, May 2000.

[7] J. Fridrich and Miroslav Goljan. Robust hash functions for digital watermarking. In The International Conference on Information Technology: Coding and Computing, pages 178 –183, 2000.

[8] Celia Gorman. Counterfeit chips on the rise. IEEE Spectrum, 2012.

[9] Jaap Haitsma and Ton Kalker. A highly robust audio fingerprinting system with an efficient search strategy. Journal of New Music Research, 32(2):211– 221, 2003.

[10] V. Monga and B.L. Evans. Robust perceptual image hashing using feature points. In International Conference on Image Processing, volume 1, pages 677 – 680, October 2004.

[11] V. Monga and M.K. Mihcak. Robust image hashing via non-negative matrix factorizations. In IEEE International Conference on Acoustics, Speech and Signal Processing, volume 2, page II, May 2006.

[12] Pierre Moulin. Statistical modeling and analysis of content identification.

In Information Theory Applications, San Diego, USA, 2010.

[13] Committee on Armed Services. Inquiry into counterfeit electronic parts in the department of defense supply chain. Technical report, United States Senate, 2012.

[14] J. Picard. Digital authentication with copy-detection patterns. In R. L. van Renesse, editor, Optical Security and Counterfeit Deterrence Techniques, Proceedings of the SPIE, volume 5310, pages 176–183, Jun 2004.

[15] M. Schneider and Shih-Fu Chang. A robust content based digital signature for image authentication. In International Conference on Image Processing, volume 3, pages 227 – 230, September 1996.

[16] Claude E. Shannon. A mathematical theory of communication. Bell System

Technical Journal, 27:379–423, 1948.

[17] R. A. Steenblik and M. J. Hurt. Unison micro-optic security film. In R. L.

van Renesse, editor, Optical Security and Counterfeit Deterrence Techniques V. Proceedings of the SPIE, volume 5310, pages 321–327, Jun 2004.

[18] Ashwin Swaminathan, Yinian Mao, and Min Wu. Image hashing resilient to geometric and filtering operations. In Workshop on Multimedia Signal Processing, pages 355– 358, sept.-1 oct. 2004.

[19] Harold Henry Trimm. Forensics the Easy Way, page 276. Barron's Educa-

tional Series. Barron, 2005.

[20] Sviatoslav Voloshynovskiy, Oleksiy Koval, Fokko Beekhof, Farzad Farhadzadeh, and Taras Holotyak. Information-theoretical analysis of private content identification. In IEEE Information Theory Workshop, Dublin, Ireland, August 30 – September 3 2010.

[21] Ye Wang, Shantanu Rane, Stark C. Draper, and Prakash Ishwar. An information-theoretic analysis of revocability and reusability in secure biometrics. In Workshop on Information Theory and its Applications, San Diego, CA, February 2011.

[22] Frans M.J. Willems. Information theory and biometrics. Keynote Lecture at the Sixth International Conference on Intelligent Information Hiding and Multimedia Signal Processing, October 15–17 2010.

[23] Jacob Ziv and Abraham Lempel. A universal algorithm for sequential data compression. IEEE Transactions on Information Theory, 23(3):337–343, May 1977.

[24] Jacob Ziv and Abraham Lempel. Compression of individual sequences via variable-rate coding. IEEE Transactions on Information Theory, 24(5):530– 536, September 1978.