Multimedia Authenticity with ICA Watermarks

Steven Noel[a] and Harold Szu[b]

[a]University of Louisiana at Lafayette, Lafayette LA 70504
[b]Office of Naval Research, Code 313, Arlington VA 22217-5660

ABSTRACT

For the first time the blind source de-mixing is applied to authenticity protection for multimedia products. We give an overview of the current state of multimedia authenticity protection, including the requirements of various multimedia applications, current approaches to the problem, and the robustness of the approaches. We then introduce the de-mixing algorithm, based on independent component analysis (ICA) seeking statistically factorized probability density and yielding a fast de-mixing computing using unsupervised artificial neural networks. We describe how their blind demixing capability extends signal processing from the conventional one-sensor approach to a multi-sensor approach, as in the 2 eyes and 2 ears of human sensor systems but packaged in a spatio-temporal multiplexing fashion. For trademark security, a covert ICA can serve as a dormant digital watermark embedded within the multimedia data. Unauthorized removal of the trademark as plagiarism could degrade the quality of the content data (not shown here). We show how these new approaches contribute to a flexible, robust, and relative secure system for protecting the authenticity of multimedia products.

Keywords: Multimedia watermark, copyright protection, e-commerce, independent component analysis, unsupervised neural networks.

1. INTRODUCTION

While the Internet provides ready access to products such as multimedia, it also makes unauthorized copying relatively easy. The need to protect against unauthorized copying is certainly not a new problem. One tool for verifying authenticity in the traditional printed medium is the watermark, which has begun to be extended to digital multimedia [1][2][3]. However, the richness of multimedia offers much more flexibility for protecting authenticity than mere watermarking. For example, authenticity marking could be done in some other medium than the original product, or the marking could be distributed over multiple media. Marking could be either overt, covert, or both. This opens the opportunity for authenticity marking to serve more general business purposes, such as marketing and advertising.

Inherent in multimedia watermarking is the need to mix host signals such as images, audio, or video with one or more marking images/signals. At some point the host and marking signals must then be demixed, typically by an authorized person attempting to prove authenticity. One could rely on the multimedia watermark decoder having available the original unmixed host and marking signals, as well as precise knowledge of the mixing process. But, that would require the expensive overhead of saving the original signals in a central repository, as well as render the customers' decoders obsolete if the form of the mixing were changed for any reason.

A much more flexible system would be possible if the multimedia decoder could demix without having the original host and mark signals, or without knowing exactly how they were mixed. This is known as blind source mixing. With blind demixing, a variety of multimedia signals could be mixed into a single convenient multicast data stream, with the signals being demixed by the paid customer.

Recently, excellent progress has been made toward blind demixing via unsupervised neural networks [4], an approach known as independent component analysis (ICA) [5][6]. The original signals are modeled as statistically independent signal components, which are then linearly mixed. Here independence is over all orders of statistics, not just

Further author information –
S.N.: Email: sen7849@louisiana.edu; Fax: 318-482-5791
H.S.: Email: szuh@onr.navy.mil; Telephone: 703-696-6567; Fax: 703-696-1331

2nd-order correlations. Neural networks with unsupervised learning rules are able to estimate the signal mixing to high accuracy, so that the original signals (independent components) can be recovered through inverse mixing.

A novel aspect of ICA neural networks is that they extend conventional single-sensor processing of signals to multiple sensors, as in human sensor systems (e.g. eyes and ears). The networks are able to simultaneously compare sensor outputs, extracting noise so that only coherent signals remain, based on the principle that what is not noise must be signals. The important lesson is that we must go beyond the least-mean-square error energy criteria of traditional single-sensor signal processing. Indeed, multimedia computers of the future may be equipped with ICA unsupervised learning chips for sensory preprocessing.

The ability to blindly demix signals enables a novel form of security against those who may attack multimedia watermarks. One of the independent component signals in the mixed data stream could be a copy of an overt (visible or audible) authenticity mark on the host signal. A multimedia player enabled with an ICA neural network could then blindly demix the overt mark copy and determine whether the mark is still present in the host data. If the overt authenticity mark has been removed, the player could then release a process that degrades the quality of the host data with a certain level of noise. The necessity of a demixing encryption key for authorized users would ensure that a trusted multimedia player is used. An additional level of security an d convenience could be provided by object-based technologies such as the platform-independent Java language.

2. MULTIMEDIA AUTHENTICITY PROTECTION FOR INTERNET COMMERCE

The potential for commercial gain on the Internet is vast and growing. The amount they are spending is doubling almost twice that fast. The lack of physical constraints allows a snowball effect for the popularity of Internet multimedia products. The distribution of free product samples is the gravity for snowballing success. However, it is critical that proper ownership credit remains with the product. In the case of scholarly or artistic works, allowing copying is critical for exposure, as long as the author or artist is cited. The bottom line is that producers want their products to become popular, but do not want others to copy in order to steal their market share.

The digital watermark is a first-step method for protection against unauthorized copying and deletion of the media trademark for stealing or infringing upon the host multimedia data. But a mere watermark is not sufficient to protect plagiarism. Much more elaborate methods of counter-plagiarism have been developed, but are beyond the scope of this paper. Watermarking is taken from the field of steganography, which is related to cryptography. While the goal in cryptography is to hide the meaning of covert data, the goal in steganography is to hide even the presence of it. In this case, hidden data means that it is perceptually and statistically undetectable. The true lack of perceptibility means the hidden data does not degrade the multimedia product.

Embedded data should be able to survive routine signal processing of the host data, for example by image processing programs such as Premier and Adobe PhotoShop. Ideally, watermarks should be robust with respect to operations such as filtering, cropping, rescaling, resampling, and image shear. An operation that causes particular trouble for most watermarking schemes is image rotation. Programs exist for removing embedded image data [8][9]. They apply local block-based geometric distortions (stretching, shearing, rotation), followed by resampling using bilinear interpolation.

Watermarking systems should also be able to survive the conversion to analog, then back to digital. Examples include intermediate copying to VHS tape, or digital scanning of images that have been printed to paper. Watermarks should also remain after lossy compression of the host data. While multimedia products may be attacked by the deliberate addition of noise, for most schemes the necessary level of added noise renders the host data completely unusable. Another possible attack is to average over multiple multimedia products that all have the same watermark, improving the chance of detecting the hidden data.

Some schemes allow the embedding of multiple watermarks, which is necessary for some applications like generational control (e.g. one-copy-only video). However, most of these schemes are unable to resolve the deadlock that occurs when multiple claims to ownership are made. In other words, someone could attack a multimedia product by placing his or her own mark on it, in addition to the rightful one. It has been proposed to use a trusted 3rd party timestamping service to resolve deadlocks [10]. However, it would still be possible for someone to get a timestamp in the present, then use it in the future to make an illegitimate mark. Moreover, it would be infeasible for the timestamp service to keep records of the content of multimedia products for which is it has provided timestamps.

It is possible to obscure the meaning of the hidden data by encrypting it before embedding it in the host data. That is, the information is scrambled by an encryption transformation, then sent in its scrambled form. After receipt, it is unscrambled by the inverse transformation. Effective encryption relies on the fact that no one has found a way to factor composites of 2 very large prime numbers (on the order of 200 digits) within a reasonable time. However, encryption alone merely restricts access to data, and does not label ownership. Once the data is decrypted, the signature is removed and no proof of ownership remains. It is possible to use either private or public decryption keys. But public-key watermarking would be more vulnerable, since it is much easier to remove embedded data that is already known. There have been some initial results at resolving this dilemma, by using a combination of both public and private keys [11].

Shannon gave a theoretical guarantee for the security of data, regardless of the computational power of the opponent [12]. Simmons has developed a corresponding guarantee for the authentication of data [13]. While somewhat general models of data hiding have begun to be developed [14][15][16], there current exists no theory that guarantees the hiding of data regardless of the computational power of the opponent. Thus, so far no watermarking scheme can give unconditional covertness. Progress in this area has been hampered by fact that classical intuitions in information theory often serve poorly in steganography. But one initial result is that data hiding is generally only possible where compression is inefficient to some extent, to provide some degree of redundancy for embedding hidden data [14].

The earliest watermarking schemes operated directly in the spatial or temporal domains. The simplest strategy is to encode hidden data in the least-significant bits of the host data. However, in this approach the hidden data is extremely easy to detect and remove. Most current methods are forms of spread spectrum communications, as popularized by Cox *et al.* [17][18]. They are inherently a projection on a given key-defined direction. The general idea is that the sender and receiver share a secret key for generating a pseudo-random stream of numbers, a so-called keystream. The keystream is then used to project the hidden data onto the host data.

One strategy is for the keystream to select the samples in which to embed the hidden data, thus frustrating detection [19]. But this is still vulnerable to simple filtering, which modifies the value of many of the least-significant bits. Countermeasures to this are using error correcting codes, or relying on redundancy. The extreme form of redundancy is to simply add the keystream itself into the host data, then later detect its presence via matched filter correlation [17]. While in this method the embedded data is indistinguishable from noise, only a single bit of information is embedded.

More bits can be embedded with a less redundant approach. Here the hidden bits modulate the keystream, for example by modifying its sign. Redundancy is retained by modifying multiple keystream values for each embedded bit. Overall, spread-spectrum techniques in the spatial or temporal domains have low hidden-bit capacities, e.g. 8 bits per frame for video [11], or 42 bits per second for audio [20]. They are also particularly vulnerable to attacks that break up synchronization of the keystream, such as cropping, line dropping, or random deletion and duplication of samples.

A fundamental problem in steganography is that the bits in which one can safely embed hidden data are by definition redundant, or at least nearly so. After all, an attacker would be unaware of their alteration. These redundant bits may well be removed by efficient lossy compression. Thus, there is obviously a close interaction between data hiding and data compression. If we know in advance what type of compression scheme will be used, it can guide our approach to data hiding, and greatly improve hidden data capacity.

This is the motivation for performing watermarking in the compression transform domain. Essentially every lossy compression scheme transforms the data into a sparse representation, e.g. block DCT, wavelet, or fractal. It is in the transform domain that lossy compression occurs, in the form of quantization. While the quantized transform coefficients are further compressed by entropy coding, this step is lossless and therefore has no effect on data hiding. If the data is already compressed, embedding watermarks in the compression transform domain is also more efficient, since expensive decompression is avoided. For maximum robustness, the watermark should be embedded into the same transform components already populated by the host data.

One of the more robust and popular transform-based spread spectrum schemes for image watermarking is found in [18]. They apply the 2-dimensional DCT to the entire image, select the 1000 most significant DCT components, and then add a random Gaussian keystream to these components. A comparison of block-DCT and wavelet approaches is given in [21]. The conclusion is that the block-DCT approach has the advantage of being able to watermark partially decompressed JPEG images, while the wavelet approach is much more robust. See [2] for an in-depth review of current watermarking schemes.

The fact that humans consume multimedia signals implies that the embedding of hidden data is possible. Human sensors (e.g. eyes and ears) are imperfect detectors. An audio or video signal must have a minimum intensity level before a human can detect it. However, not every sample of the multimedia signal is necessarily suitable for embedding hidden data. Whether or not a signal may be perceived depends on the spatial, temporal, and frequency characteristics of human sensor systems.

This motivates the design of perceptual masks, adapted to human sensors as well as the host data. For example, humans are more sensitive to image variations in flat areas, and less sensitive to variations in edges. For audio, when 2 tones close in frequency are played together, the louder one will mask the quieter one. Also, when a loud sound is heard, it takes time before a quiet one can be heard. Signals that are slightly perceptible may become imperceptible in the presence of masking signals. Masking models used in data hiding were often originally developed for data compression [20][21].

There are numerous applications of data embedding specific to video. One is the tracking of commercials for automatic billing (pay per play). Also, systems that allow multiple watermarks could be used for generational control of video, for example one-copy-only DVD [22]. If fact, both DVD [23] and MPEG-4 [24] require that watermarks be decoded by any receiver, with no possibility of faking the marks. Other data embedding applications for video include multicasting of movies in different languages and ratings.

For audio, the most important application of watermarking is ownership protection in music [25]. Internet sites distributing pirated music compressed in MPEG audio layer 3 are relatively common. The Recording Industry Association of America (RIAA) and the International Federation of the Phonographic Industry (IFPI) have proposed requirements for the robustness of embedded audio data against signal processing and deliberate attacks. Their 1997 Request for Proposals was evaluated by the MUSE Project [26], jointly funded by the recording industry and the European Union. The test results were somewhat weak with respect to watermark audibility and robustness, leading to a 2nd round of proposals.

Another potentially important application of the hidden marking of multimedia data is in medical imagery. Current medical formats separate image data from its corresponding descriptive text such as patient name, date, and physician. Thus, the link between image and patient can occasionally be broken. Embedding the patient's name directly in the image data provides an additional measure of safety.

3. BLIND DEMIXING OF MULTIMEDIA DATA WITH ICA NEURAL NETWORKS

Many watermarking systems are termed non-blind, in the sense that they require original versions of both the host data and watermark keystream in order to extract the mark once it is embedded. This may be a reasonable requirement for applications that merely verify the existence of the watermark for proving ownership. The owner could be expected to have access to the original data. However, non-blind schemes are infeasible for systems in which the *consumer* rather than the producer needs to extract useful embedded data. This would allow valuable features to be added to the multimedia product, rather than just mere copyright protection.

There do exist watermarking schemes in which access to the original host data is not required. While these are sometimes called blind schemes, we consider them only semi-blind, since the watermark keystream is still required. The hidden data capacity of these semi-blind systems is limited by the reduced dynamic range of the relatively weak embedded signal, as well as by the redundancy necessary for making this weak signal robust.

Some applications of data embedding do indeed require the embedded data to be hidden, such as an undetectable ownership mark. However, hiding the embedded data is not always the best solution. For example, an unhidden visible or audible mark could claim ownership overtly. Such a mark leaves no question as to product ownership, and provides the product exposure that is vital for snowballing popularity.

Indeed, there are a multitude of applications in which the consumer should have access to the embedded data. Examples include video multicasting or medical image identification. In short, allowing embedded data to be overt as well as covert can promote visibility, flexibility, and creativity in addition to providing mere authentication.

Relaxing the requirement of imperceptibility for embedded data removes the need for redundant coding for embedded signal robustness. This lack of redundancy in turn increases the bitrate of the embedded data, to the full rate of the host signal. Allowing embedded signals to be perceptible eliminates the need for perceptual masking models. It also removes the requirement that the embedded signals be statistically undetectable.

But this raises problems. If signals are embedded at the full strength of the host signal, the line between host and embedded signals blurs. Also, since the embedded signals are now perceptible, they interfere with the host signals, affecting their quality. How then are host and embedded signals to be separated?

A straightforward model for combining host and embedded signals is to form a simple linear combination of them. The signals could then later be demixed given the linear mixing coefficients. However, this provides little in the way of flexibility or security. It would be much better if the signals could be demixed without knowledge of the original mixing coefficients. Our scheme would be even more flexible if we could demix without knowledge of the original signals themselves. Thus we are faced with the problem of truly blind demixing of source signals.

Fortunately, recently developed neural networks with unsupervised learning have accomplished blind demixing of signal [4]. The approach is known as independent component analysis (ICA) [5][6]. The source signals are modeled as statistically independent signal components, which have been subsequently linearly mixed. Independence is defined such that the joint probability densities of the signal components can be factorized as the product of the marginal densities. Thus independence is over all orders of statistics, not just 2nd-order correlations.

The neural networks with unsupervised learning are able to estimate the signal mixing to high accuracy, so that the original signals (independent components) can be recovered through inverse mixing. The learning rules are based on maximizing the degree to which the independent components are nongaussian, which is justified by an appeal to the Central Limit Theorem. Possible measures of nongaussianity are the absolute value of kurtosis (4th-order cumulant), or negentropy (slightly modified version of differential entropy).

The ICA model assumes that there are multiple sensors, each sensing a different mix of the independent components. Thus is extends conventional single-sensor processing of signals to multiple sensors, analogous to the multiple sensors (eyes and ears) in humans. The unsupervised ICA neural networks are able to simultaneously compare sensor outputs, extracting noise so that only coherent signals remain. In the unsupervised learning rule there is no specific desired output other than white gaussian noise. This is consistent with the idea that what is not noise must be signals.

Figure 1 is an example of the mixing of multimedia data, with blind demixing by an ICA unsupervised neural network. The 3 independent data components are the host data, an advertisement, and an authenticity mark. The data are mixed to insure the various components remain together for Internet distribution. The neural network with unsupervised learning estimates with high accuracy the mixing matrix **A**, allowing blind demixing of the independent components (source images).

4. INDEPENDENT COMPONENTS FOR MULTIMEDIA WATERMARK

For Internet multimedia products, the marketing strategy is typically to make product samples freely available, then make the products themselves available only after a fee is paid. We can learn from the example of the radio music industry, in which the more free exposure is given, the more popular the product becomes. A strategy that is often successful on the Internet is to provide a product for free, but to charge for service of the product. An example is Adobe Acrobat, in which the reader software is free, but the software to actually write documents must be purchased.

The dilemma for multimedia producers is controlling access to the products that are not distributed for free. Watermarking may be helpful in this regard. However, most watermarking efforts have been aimed at embedding data that is hidden. Figure 2 shows such a standard model for watermarking of multimedia products. The multimedia data is playable via standard media players. Only the host multimedia data is visible, but a hidden watermark exists for future copyright proof. As we argued in the previous section, it is often advantageous for the embedded data to be unhidden, as in overt perceptible ownership marks or advertisements. This leads to the need for blind demixing of the source signals, which has been done to high accuracy with ICA unsupervised neural networks.

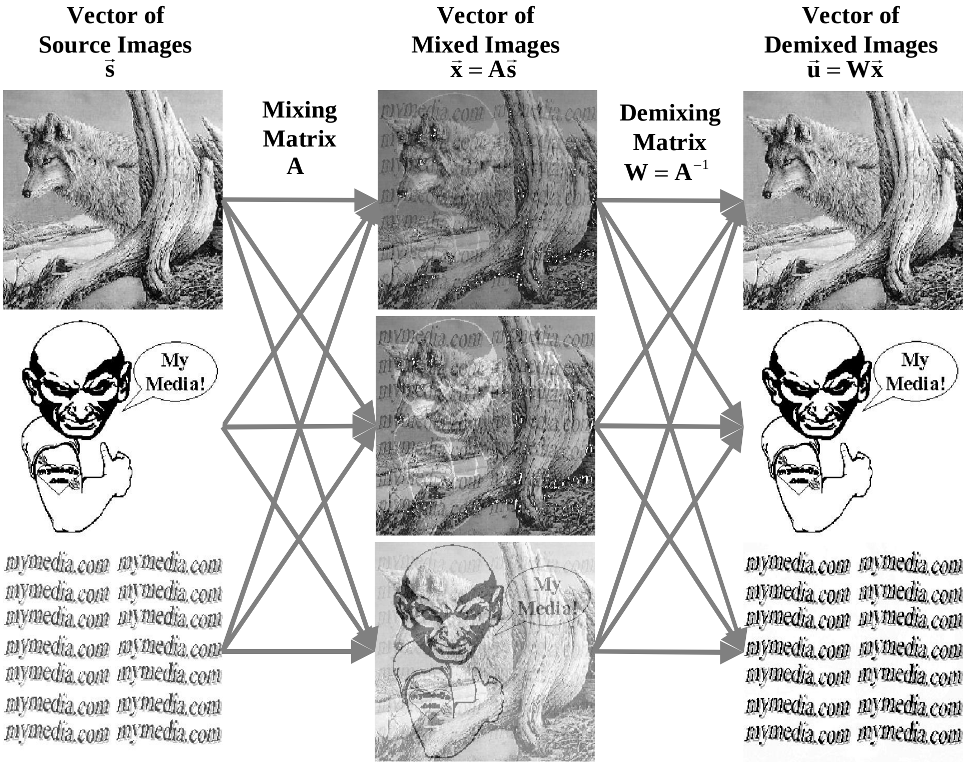
Vector of Source Images $\vec{s}$
Mixing Matrix A
Vector of Mixed Images $\vec{x} = A\vec{s}$
Demixing Matrix $W = A^{-1}$
Vector of Demixed Images $\vec{u} = W\vec{x}$
My Media!
mymedia.com
My Media!
My Media!
mymedia.com

Figure 1. Example mixing of multimedia data, with blind demixing via ICA unsupervised neural network

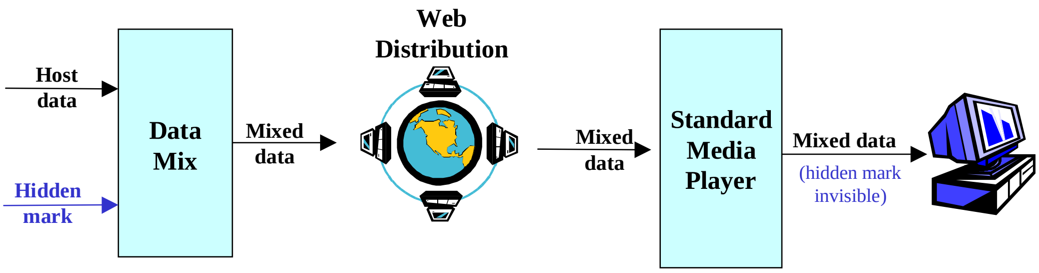

Figure 2. Standard model for watermarking of multimedia product

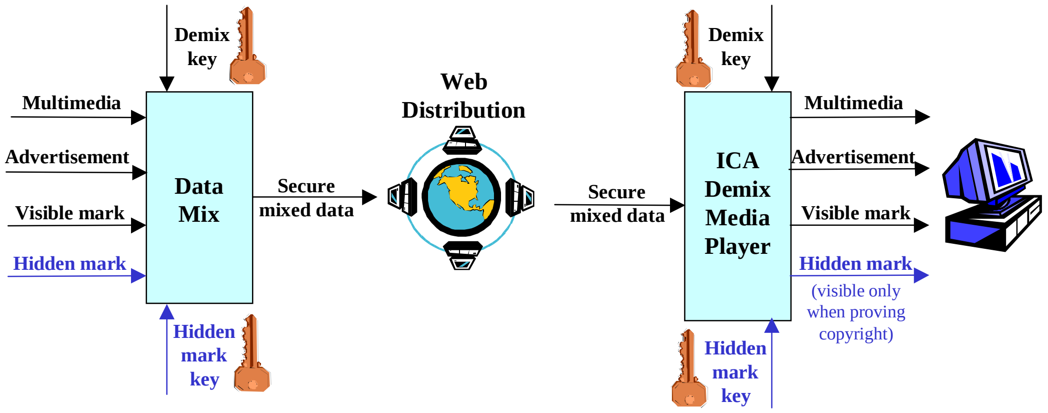

Figure 3. Model for neural net ICA demix enabled media player

As shown in Figure 3, a visible mark and an accompanying advertisement are possible through neural net ICA demix enabled media players. The demix key allows the data to be demixed, providing access to original multimedia data, a visible mark, and an advertisement. The hidden mark key allows the hidden mark to be demixed for copyright proof (the demix key alone is insufficient for accessing the hidden mark). The demix key can be omitted to allow free public access to all data except the hidden mark. It is important that ICA demixing be robust with respect to various signal processing. Being linear, ICA demixing is invariant with respect to linear filtering. It is also invariant with respect to various changes in sampling which cause synchronization problems with many other data embedding schemes. Such sampling changes include cropping, line dropping, or changing the sample rate. In the ICA model, each sample point is a mix of 2 independent components, and the mix is same for all samples.

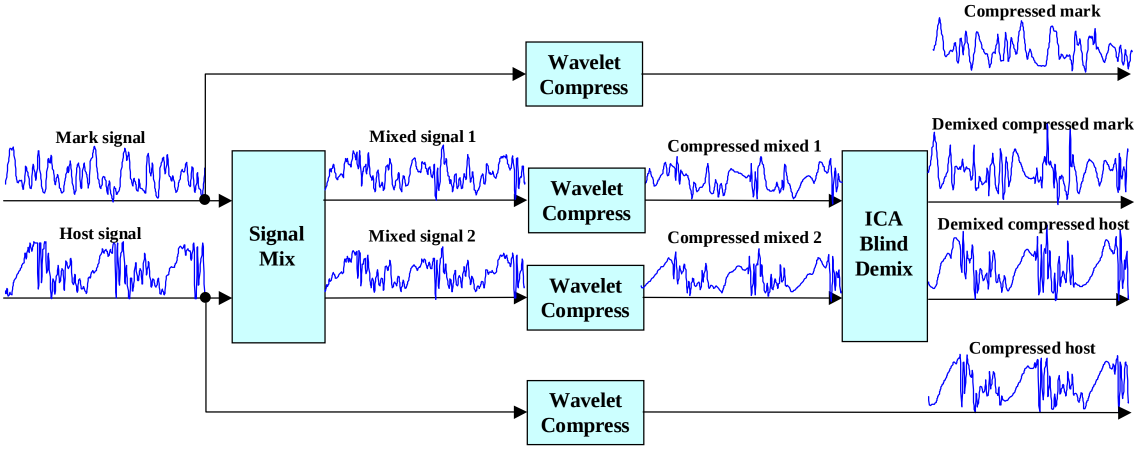

Figure 4. Robustness of ICA neural net blind demixing with respect to nonlinear wavelet audio data compression

It is also critical that ICA neural networks be able to demix signals that have been subjected to lossy compression. In general, compression can occur either before or after mixing. Compression before mixing yields nearly lossless demixing, and avoids costly decompression of previously compressed data. As we see in Figure 4 for audio data, ICA neural networks can blindly demix signals with minimal distortion even after they have been subjected to the nonlinear transformation of lossy wavelet compression.

In Figure 5, after the host and watermark images (a) and (b) are mixed, the mixtures are subjected to DWT compression. ICA unsupervised neural nets are still able to blindly demix host and mark images, as shown in (e) and (f), despite the nonlinear transformation of mixed images via compression. Compare these with images (c) and (d), which are simply compressed versions of the original images, with no mixing involved. Note that this is not an issue if the host and mark images are compressed separately before mixing them, in which case the demixed images are nearly indistinguishable from the originals. Mixed signals must also be protected against piracy. This is particularly important for ICA blind demixing, in which multiple sensor looks are involved. We propose the protection scheme in Figure 6. The mixing of host and mark signals is randomized, within random-width local windows. An authorized user needs a key to demix the mixed signals. This is easily generalized to images by using random size blocks. This is consistent with prudent cryptographic practice, as first described by Kerckhoff in1883. One must assume that the method to encipher data is known to an opponent, since the method may be compromised, thereby compromising all protected data. Security lies in the choice of key, in this case the random number seed.

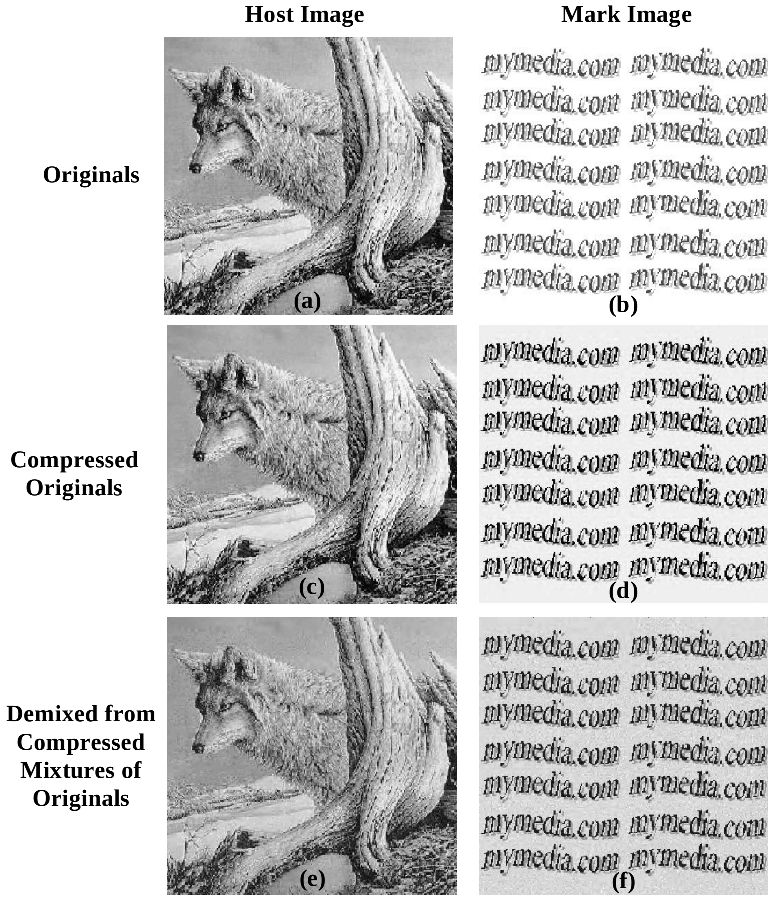

Figure 5. Robustness of ICA neural net blind demixing with respect to nonlinear wavelet image data compression

This scheme protects the mixed data from being demixed by a pirate, so that only an authorized customer is able to demix it. But this is not secure enough. Once the data is demixed, it is open to attack by the authorized customer himself. For example, the customer could attempt to remove a visible overt company logo from the demixed multimedia product, to claim ownership. We do not have the luxury of a DVD-like approach, where there is industry-wide cooperation and all multimedia players are trusted. Our approach needs to be completely autonomous.

We also propose a novel method for degrading the quality of multimedia data when it is attacked. We avoid the need to compute pseudo-random noise for degrading the data. We instead degrade it by simple manipulation of the wavelet transform coefficients in the compressed domain, as shown in Figure 7. Fast degradation of original image (a) is possible by simply shifting the lowpass DWT coefficients by a certain number of samples, e.g. shifting by 2 samples in (b) or by 5 samples in (c). Stronger degradation is possible by reversing decoding order of lowpass coefficient order within groups, e.g. reversing in sets of 2 coefficients in (d) or reversing in sets of 5 coefficients in (e).

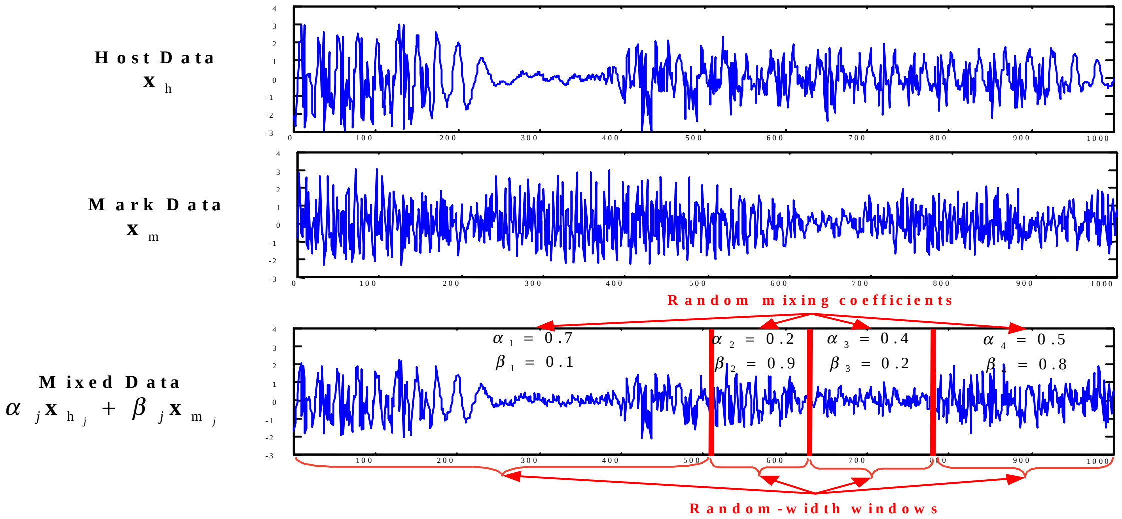

Figure 6. Localized random mixing for securing mixed signals against piracy

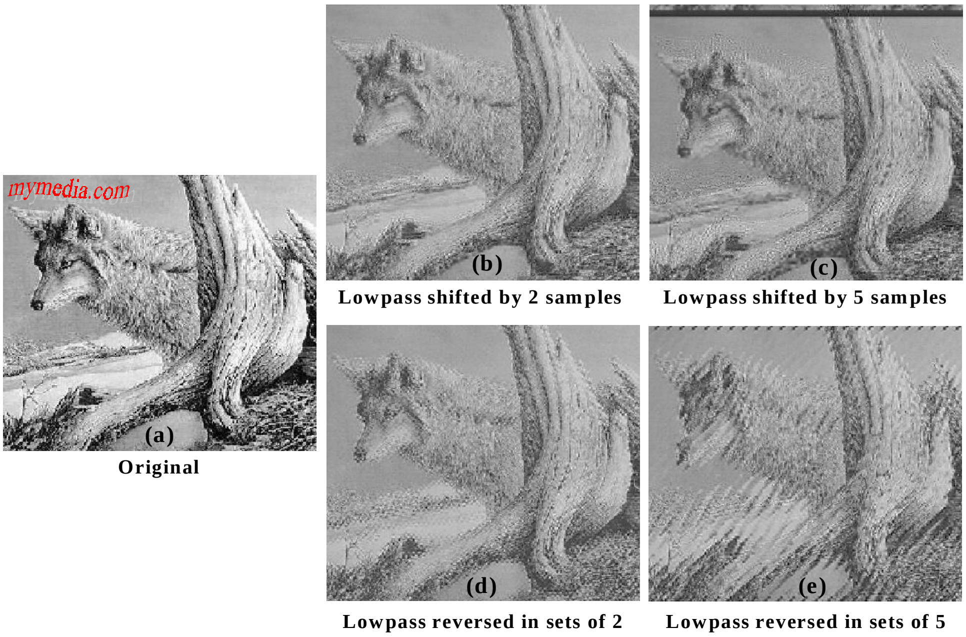

Original

Lowpass shifted by 2 samples

Lowpass shifted by 5 samples

Lowpass reversed in sets of 2

Lowpass reversed in sets of 5

Figure 7. Fast degrading of image quality by simple changes to decoding of DWT coefficients

We recognize a way in which to implement our scheme of watermark. Object technology such as the platform-independent Java language could provide a particularly convenient implementation. Multimedia data, advertisement, and mark data are mixed and secured, then embedded within the object. The neural network ICA demix media player is defined as an object method, with read-only access to mixed data. The demix key allows data to be demixed, and the mark key allows the hidden mark to be accessed for copyright proof.

5. SUMMARY AND CONCLUSIONS

We have examined the application of independent component analysis (ICA), via unsupervised neural networks, to authenticity protection for multimedia products. We described how the blind demixing capability of these networks extends signal processing from a one-sensor approach to a multi-sensor approach. Unauthorized removal of the watermark degrades quality of the product being pirated. We showed how this approach could provide a flexible, robust, and secure system for protecting the authenticity of multimedia products.

REFERENCES

1. R. Barnett, "Digital Watermarking: Applications, Techniques, and Challenges," *Electronics and Communication Engineering Journal*, 11(4), 173-183, August 1999.
2. F. Hartung, M. Kutter, "Multimedia Watermarking Techniques," *Proceedings of the IEEE*, 87(7), 1079-1107, July 1999.
3. M. Swanson, M. Kobayashi, A. Tewfick, "Multimedia Data-Embedding and Watermarking Technologies," *Proceedings of the IEEE*, 86(6), June 1998.
4. H. Szu, "Progress in Unsupervised Artificial Neural Networks for Image Demixing Applications," *IEEE Industrial Electronics Society Newsletter*, 46(2), 7-12, June 1999.
5. P. Comon, "Independent Component Analysis – A New Concept?," *Signal Processing*, 36, 287-314, 1994.
6. A. Hyvarinen, E. Oja, *Independent Component Analysis: A Tutorial*, tutorial notes for International Joint Conference on Neural Networks (IJCNN'99), Washington D.C., http://www.cis.hut.fi/projects/ica/, July 1999.
7. "Getting Started with Electronic Commerce," *The Journal*, January 26, 2000.
8. UnZign software at http://altern.org/watermark.
9. M. Kuhn, *Stirmark*, http://www.cl.cam.ac.uk/~fapp2/watermarking/image_watermarking/stirmark, 1997.
10. Digital timestamping service found at http://www.surety.com.
11. F. Hartung, B. Girod, "Fast Public-Key Watermarking of Compressed Video," *International Conference on Image Processing*, Volume 1, 528-531, 1997.
12. C. Shannon, "Communication Theory of Secrecy Systems," *Bell Systems Technical Journal*, 28, 656-715, 1949.
13. G. Simmons, "A Survey of Information Authentication," in *Contemporary Cryptology – The Science of Information Integrity*, IEEE Press, New York, 379-419, 1992.
14. R. Anderson, F. Petitcolas, "On the Limits of Steganography," *IEEE Journal on Selected Areas in Communications*, issue on Copyright and Privacy Protection, 16(4), 474-481, May 1998.
15. J. Hernandez, F. Perez-Gonzales, J. Rodriguez, G. Nieto, "Performance Analysis of a 2-D-Multiphase Amplitude Modulation Scheme for Data Hiding and Watermarking of Still Images," *IEEE Journal on Selected Areas in Communications*, issue on Copyright and Privacy Protection, 16(4), 510-524, May 1998.
16. B. Chen, G. Wornell, "An Information-Theoretic Approach to the Design of Digital Watermarking Systems," *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 1999.
17. I. Cox, J. Killian, T. Leighton, T. Shamoon, "Secure Spread Spectrum Watermarking for Multimedia," NEC Research Institute, Technical Report 95-10, 1995.
18. I. Cox, J. Killian, T. Leighton, T. Shamoon, "Secure Spread Spectrum Watermarking for Images, Audio, and Video," *IEEE International Conference on Image Processing*, Volume 3, 243-246, 1996.
19. E. Franz, A. Jerichow, S. M🌧ller, A. Pfitzmann, I. Stierand, "Computer Based Steganography," in *Information Hiding, Springer Lecture Notes in Computer Science*, Volume 1174, 7-21, 1996.
20. M. Swanson, B. Zhu, Z. Tewfick, "Current State of the Art, Challenges, and Future Directions for Audio Watermarking," *IEEE International Conference on Multimedia Computing and Systems*, Volume 1, 19-24, 1999.
21. C. Podilchuk, W. Zeng, "Image-Adaptive Watermarking Using Visual Models," *IEEE Journal on Selected Areas in Communications*, issue on Copyright and Privacy Protection, 16(4), 525-538, May 1998.
22. Data Hiding Subgroup (DHSG) of Copy Protection Working Group (CPWG), http://www.dvcc.com/dhsg, May 1997.
23. Watermarking for DVD – Call for Proposals, http://www.dvcc.com/dhsg/, July 1997.
24. MPEG-4 Requirements Group, Call for Proposals for Identification and Protection of Content in MPEG-4, ICO document JTC1/SC29/WG11 N1714, April 1997.
25. P. Jessop, "The Business Case for Audio Watermarking," *IEEE International Conference on Acoustics, Speech, and Signal Processing*, 2077-2074, 1999.
26. IFPI MUSE Project: Embedded Signaling, http://www.ifpi.org/technology/muse_embed.html.