

Universal Decoding of Watermarks Under Geometric Attacks

Pierre Moulin

University of Illinois

Beckman Inst., Coord. Sci. Lab & ECE Dept.

405 N. Mathews Ave., Urbana, IL 61801

Email: moulin@ifp.uiuc.edu

Abstract—Designing watermarking codes that can withstand geometric and other desynchronization attacks is a notoriously difficult problem. One may ask whether these difficulties are due to limitations of current codes, or rather to fundamental limitations on achievable performance. We model the attack channel as the cascade of a memoryless channel and a smooth, invertible mapping $T_{\theta}, \theta \in \Theta_n$, representing the geometric attack. The decoder does not know the value of θ . We show that under regularity conditions, there exists a universal decoder for this problem, and we explicitly identify it.

I. INTRODUCTION

One of the main difficulties in designing watermarking and data-hiding codes is to ensure a certain level of robustness against desynchronization and other geometric attacks. Such attacks include image warping, amplitude modulation, and for audio and video signals, temporal desynchronization [1]–[4]. If the original host signal is available to the decoder (*nonblind watermarking*), there is clearly hope to “undo” these attacks with the help of this signal, hence an approach based on joint signal registration and message decoding is plausible. In the opposite situation where the host is not available to the decoder (*blind watermarking*), the task seems much harder. One may wonder whether this increased difficulty is due to the deficiencies of current code designs, or to some fundamental performance limit.

This paper describes our recent results in that direction, starting from a mathematically tractable formulation of the problem. The attack is modeled as the cascade of a memoryless channel and a smooth, invertible mapping representing the geometric attack. This mapping is parameterized by an unknown parameter θ (e.g., a scaling parameter, a filter response, or a time-warping function). We address the potential loss in error probability due to lack of knowledge of θ by the receiver.

The framework for this study is universal decoding [5]. The noncoherent penalty (due to lack of knowledge of θ by the decoder) is measured by the loss of random-coding error exponents relative to the coherent case (in which the decoder knows θ). A universal decoder incurs no loss in random-coding exponents. For many problems, universal decoders do not exist. Furthermore, even when universal decoders exist, their mathematical structure could be prohibitively complex (e.g., the merged list decoders of [5].) Obvious choices such as Generalized Maximum Likelihood decoders (which jointly estimate θ and decode the transmitted message) are generally not universal.

In the watermarking problem, the invertibility of the family of geometric mapping plays a key role in the analysis. Provided that family has limited “complexity” (in a sense made precise below), we show that universal decoders exist and admit a relatively simple structure. Detailed derivations and proofs are given in [6].

II. NOTATION

We use uppercase letters to denote random variables, lowercase letters to denote their individual values, and boldface letters to denote sequences of symbols defined over a common alphabet, e.g., $\mathbf{x} = (x_1, \dots, x_n) \in \mathcal{X}^n$. Given a sequence $\mathbf{x} \in \mathcal{X}^n$, we denote by $p_{\mathbf{x}}$ its type (a pmf defined over the alphabet \mathcal{X}). Given two sequences $\mathbf{x} \in \mathcal{X}^n$ and $\mathbf{y} \in \mathcal{Y}^n$, we denote by $p_{\mathbf{xy}}$ their joint type (a pmf over $\mathcal{X} \times \mathcal{Y}$), and by $p_{\mathbf{y}|\mathbf{x}}$ the conditional type of \mathbf{y} given \mathbf{x} . Finally, let $d_H(\mathbf{x}, \mathbf{y}) = \#\{i : x_i \neq y_i, 1 \leq i \leq n\}$ denote the Hamming distance between \mathbf{x} and \mathbf{y} . The notation $f_n \doteq g_n$ denotes asymptotic equality on the exponential scale, i.e., $\lim_{n \rightarrow \infty} \frac{1}{n} \log(f_n/g_n) = 0$. The symbol \mathbb{E}_X denotes expectation with respect to a random variable X .

III. MATHEMATICAL MODEL

We adopt the following communication model for watermarking, see Fig. 1 [7]. Given a host sequence $\mathbf{s} \in \mathcal{S}^n$,

side information $\mathbf{k} \in \mathcal{K}^n$, and a message m uniformly distributed over $\{1, \dots, \lceil 2^{nR} \rceil\}$, the encoder produces a marked sequence $\mathbf{x} = f_n(\mathbf{s}, m, \mathbf{k})$ where f_n is the encoding function. The marked sequence \mathbf{x} is subject to attacks, resulting in a degraded sequence \mathbf{y} . The decoder returns an estimate $\hat{m} = \psi_n(\mathbf{y}, \mathbf{k})$ of the message that was sent. The side information \mathbf{k} may be a cryptographic key, independent of \mathbf{S} (*public watermarking*); \mathbf{k} may also contain full information about \mathbf{S} (*private watermarking*).




Fig. 1. Communication model for watermarking and data hiding.

Referring to Fig. 2, the attack is modeled by the cascade of a fixed memoryless channel $W(z|x)$ and an invertible global mapping T_θ representing a geometric attack. Therefore $\mathbf{y} = T_\theta \mathbf{z}$, where \mathbf{z} is generated according to the pmf $W^n(\cdot|x)$. The feasible set for θ is denoted by Θ_n . The alphabets \mathcal{S} , \mathcal{X} and \mathcal{Z} are assumed to be finite.




Fig. 2. Model for geometric attacks.

The family $\{T_\theta, \theta \in \Theta_n\}$ satisfies the following condition:

- (A1) The mapping $T_\theta : \mathcal{X}^n \rightarrow \mathcal{Y}^n$ is invertible for all n and for all $\theta \in \Theta_n$.

We denote by $U_\theta = T_\theta^{-1}$ the inverse mapping. In addition, Θ_n satisfies one of the following conditions.

- (A2) The parameter set Θ_n is discrete, and its cardinality is fixed or grows subexponentially with n : $\limsup_{n \rightarrow \infty} \frac{1}{n} \ln |\Theta_n| = 0$.
- (A2') There exists a sequence $\epsilon_n \downarrow 0$ and a sequence of sets $\tilde{\Theta}_n \subseteq \Theta_n$ that satisfies the following two conditions.

- The cardinality of these subsets is upper bounded by a subexponential function of n : $\limsup_{n \rightarrow \infty} \frac{1}{n} \ln |\tilde{\Theta}_n| = 0$.

- The collection of sets $\tilde{\Theta}_n$ is dense in Θ_n in the following sense. For any $\mathbf{y} \in \mathcal{Y}^n$ and $\theta \in \Theta_n$, one can find some $\theta^* \in \tilde{\Theta}_n$ such that

$$\frac{1}{n} d_H(U_\theta \mathbf{y}, U_{\theta^*} \mathbf{y}) \leq \epsilon_n. \quad (1)$$

IV. EXAMPLES: PERMUTATIONS

Let $\theta : \{1, 2, \dots, n\} \rightarrow \{1, 2, \dots, n\}$ denote a permutation of the samples $\{1, 2, \dots, n\}$, and $T_\theta : \mathcal{X}^n \rightarrow \mathcal{Y}^n$ the corresponding permutation operator. Condition (A1) is automatically satisfied.

Example #1. If we choose Θ_n to be the set of all permutations, then Θ_n has size $|\Theta_n| = n! \doteq (n/e)^n$ (by Stirling's formula), which is superexponential in n , so condition (A2) is violated.

Example #2. To make Θ_n less complex, we could choose Θ_n to be the set of all block permutations, for blocks of size B . The size of Θ_n is now n/B

$$|\Theta_n| = (n/B)! \doteq \left(\frac{n}{eB}\right)^{n/B}$$

which is subexponential in n if and only if $\ln |\Theta_n| \ll n$, or equivalently, $B \gg \ln n$.

Example #3. Choose ϵ_n such that $\frac{1}{\ln n} \ll \epsilon_n \ll 1$, and let Θ_n be the set of permutations affecting only $n\epsilon_n$ samples. Here

$$|\Theta_n| \doteq \binom{n}{n\epsilon_n} (n\epsilon_n)!$$

which is superexponential in n because $\frac{1}{n} \ln |\Theta_n| \sim 2\epsilon_n \ln n \gg 1$. Therefore (A2) is not satisfied. However (A2') is satisfied if we choose $\tilde{\Theta}_n$ consisting of a single element, the identity operator (trivial permutation).

V. COMPOSITE HYPOTHESIS TESTING

The decoding problem is a M -ary hypothesis testing problem of the form

$$H_m : \mathbf{Y} \sim p(\mathbf{y}, \mathbf{k} | \theta, m), \quad \theta \in \Theta_n, \quad 1 \leq m \leq M. \quad (2)$$

where $M = \lceil 2^{nR} \rceil$. This is a noncoherent decoding problem because θ is unknown to the receiver. When decoding rule ψ is used, the probability of error is

$$P_e(\theta, \psi) = \frac{1}{M} \sum_{m=1}^M \Pr[\psi(\mathbf{Y}, \mathbf{K}) \neq m \mid m \text{ sent}, \theta].$$

A. Decision Rules

If θ is known to the receiver, the test that minimizes error probability is the maximum likelihood (ML) decision rule

$$\psi_{\text{ML}}(\mathbf{y}, \mathbf{k}) = \underset{1 \leq m \leq M}{\operatorname{argmax}} p(\mathbf{y}, \mathbf{k} | \theta, m).$$

Denote by $P_e^*(\theta) = P_e(\theta, \psi_{\text{ML}})$ the corresponding error probability.

If the receiver does not know θ , there exists generally no decision rule that achieves $P_e^*(\theta)$, i.e., the receiver has to pay a penalty for not knowing θ .

B. Asymptotic Optimality

We focus on universal detection rules, which (when they exist) perform as well as the ML detector on the exponential scale. More precisely, a sequence of detection rules ψ_n is said to be universal if

$$\limsup_{n \rightarrow \infty} \max_{\theta \in \Theta_n} \frac{1}{n} \ln \frac{P_e(\theta, \psi_n)}{P_e^*(\theta)} = 0. \quad (3)$$

C. Universal Decoding

In the absence of host sequence and side information ($\mathcal{S} = \mathcal{K} = \emptyset$), the watermarking model degenerates to a point-to-point communication problem. Let the input sequence \mathbf{x} to the channel be one of the $M = \lceil 2^{nR} \rceil$ elements of a codebook $\mathcal{C} = \{\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^M\}$. These codewords are generated independently according to a uniform distribution over a subset \mathcal{B}_n of \mathcal{X}^n . The encoder selects a message m and transmits the corresponding codeword \mathbf{x}^m .

Feder and Lapidoth studied the optimal-decoding problem for a general family of channels with memory $\{p_\theta(\mathbf{y}|\mathbf{x}), \theta \in \Theta\}$. Under a condition of *strong separability* on the family of channels, they proved the existence of universal decoders in the random coding sense.

The universal decoders of [5] are merged list decoders and have extremely large complexity. Generalized ML decoders have smaller complexity but are generally not universal in general [5]. The question of universality of the GML decoder for the watermarking problem is addressed in the following sections. The special structure of the family of channels assumed in Sec. III simplifies the analysis, even in the presence of side information.

Before concluding this section, observe that under our invertibility assumption (A1) on the mappings T_θ , the ML rule based on \mathbf{Y} coincides with the ML rule based on $\mathbf{Z} = U_\theta \mathbf{Y}$. The error probability of the ML decoder is therefore given by

$$\begin{aligned} P_e^*(\theta) &= \frac{1}{M} \sum_{m=1}^M \Pr[\psi_{\text{ML}}(\mathbf{Y}, \mathbf{K}) \neq m \mid m \text{ sent}, \theta] \\ &= \frac{1}{M} \sum_{m=1}^M \Pr[\psi_{\text{ML}}(\mathbf{Z}, \mathbf{K}) \neq m \mid m \text{ sent}] \end{aligned}$$

which is *independent* of θ .

VI. PRIVATE WATERMARKING – FINITE ALPHABETS

For private watermarking, $\mathbf{K} = \mathbf{S}$. Given \mathbf{s} , define \mathcal{B}_n as a fixed conditional type class $T_{\mathbf{x}^*|\mathbf{s}}$. The codebook \mathcal{C} is the collection of $M = \lceil 2^{nR} \rceil$ sequences \mathbf{x}^m , drawn independently and uniformly over \mathcal{B}_n . Hence \mathcal{C} is a conditionally constant composition code. The ML decoder is given by

$$\psi_{\text{ML}}(\mathbf{y}, \mathbf{s}) = \operatorname{argmax}_{1 \leq m \leq M} p(\mathbf{x}^m | U_\theta \mathbf{y}, \mathbf{s}). \quad (4)$$

Under Assumption (A2), the GML decoder is given by

$$\psi_{\text{GML}}(\mathbf{y}, \mathbf{s}) = \operatorname{argmax}_{1 \leq m \leq M} \max_{\theta \in \Theta_n} p(\mathbf{x}^m | U_\theta \mathbf{y}, \mathbf{s}). \quad (5)$$

An error is declared if the maximum over m is not unique. The GML and ML decoding rules coincide when Θ_n is a singleton.

Theorem 6.1: Assume the index set Θ_n satisfies Assumption (A2). Then

$$\max_{\theta \in \Theta_n} \frac{P_e(\theta, \psi_{\text{GML}})}{P_e^*(\theta)} \leq |\Theta_n| (n+1)^{|\mathcal{S}| |\mathcal{X}| (1+2|\mathcal{Z}|)}, \quad (6)$$

and therefore the GML decoder is universal.

Sketch of the Proof. Given θ , the conditional probability of error of a decoder ψ is the probability that the transmitted codeword \mathbf{x} is confused with some other codeword $\mathbf{x}' \in \mathcal{C}$. We have

$$P_e(\theta, \psi) = \mathbb{E}_{\mathbf{S}\mathbf{X}\mathbf{Z}} \Pr[\text{error} \mid \mathbf{x}, T_\theta \mathbf{z}, \mathbf{s}, \psi]. \quad (7)$$

The conditional error probability in the right side of (7) is computed as follows. Since any deterministic decoding rule ψ may be expressed as

$$\psi(\mathbf{y}, \mathbf{s}) = \operatorname{argmax}_{1 \leq m \leq M} \varphi_\psi(\mathbf{x}^m, \mathbf{y}, \mathbf{s})$$

(where φ_ψ represents a score function), we have

$$\begin{aligned} \Pr[\text{error} \mid \mathbf{x}, T_\theta \mathbf{z}, \mathbf{s}, \psi] &= 1 - \\ &= (1 - \Pr_{\mathbf{X}'}[\varphi_\psi(\mathbf{x}', T_\theta \mathbf{z}, \mathbf{s}) \geq \varphi_\psi(\mathbf{x}, T_\theta \mathbf{z}, \mathbf{s}, \psi)])^{M-1} \end{aligned} \quad (8)$$

where the probability in the right side is over \mathbf{x}' uniformly distributed over $T_{\mathbf{x}^*|\mathbf{s}}$. Now

$$\begin{aligned} \frac{P_e(\theta, \psi_{\text{GML}})}{P_e^*(\theta)} &= \frac{P_e(\theta, \psi_{\text{GML}})}{P_e(\theta, \psi_{\text{ML}})} \\ &= \frac{\mathbb{E}_{\mathbf{S}\mathbf{X}\mathbf{Z}} \Pr[\text{error} \mid \mathbf{x}, T_\theta \mathbf{z}, \mathbf{s}, \psi_{\text{GML}}]}{\mathbb{E}_{\mathbf{S}\mathbf{X}\mathbf{Z}} \Pr[\text{error} \mid \mathbf{x}, T_\theta \mathbf{z}, \mathbf{s}, \psi_{\text{ML}}]} \\ &\leq \max_{\mathbf{s}\mathbf{X}\mathbf{Z}} \frac{\Pr[\text{error} \mid \mathbf{x}, T_\theta \mathbf{z}, \mathbf{s}, \psi_{\text{GML}}]}{\Pr[\text{error} \mid \mathbf{x}, T_\theta \mathbf{z}, \mathbf{s}, \psi_{\text{ML}}]} \end{aligned} \quad (9)$$

where we have used the inequality [8, Lemma 2]

$$\frac{a_1 \cdots + a_L}{b_1 + \cdots + b_L} \leq \max_{1 \leq i \leq L} \frac{a_i}{b_i} \quad (10)$$

for any nonnegative sequences $\{a_i, 1 \leq i \leq L\}$ and $\{b_i, 1 \leq i \leq L\}$.

Now using (8) and the inequality [8, Lemma 2]

$$\frac{1 - (1 - t)^{M-1}}{1 - (1 - t')^{M-1}} \leq \max\left(1, \frac{t}{t'}\right) \quad \forall t, t' \in [0, 1] \quad (11)$$

we obtain

$$\frac{Pr[\text{error} \mid \mathbf{x}, T_\theta \mathbf{z}, \mathbf{s}, \psi_{\text{GML}}]}{Pr[\text{error} \mid \mathbf{x}, T_\theta \mathbf{z}, \mathbf{s}, \psi_{\text{ML}}]} \leq \max(1, \frac{Pr_{\mathbf{X}'}[\varphi_{\text{GML}}(\mathbf{x}', T_\theta \mathbf{z}, \mathbf{s}) \geq \varphi_{\text{GML}}(\mathbf{x}, T_\theta \mathbf{z}, \mathbf{s})]}{Pr_{\mathbf{X}'}[\varphi_{\text{ML}}(\mathbf{x}', T_\theta \mathbf{z}, \mathbf{s}) \geq \varphi_{\text{ML}}(\mathbf{x}, T_\theta \mathbf{z}, \mathbf{s})]}) \quad (12)$$

Given a decoding rule ψ , the definition of the score function φ is not unique. For the ML decoder, we find it convenient to use

$$\varphi_{\text{ML}}(\mathbf{x}, \mathbf{y}, \mathbf{s}) = \frac{1}{n} \ln \frac{p(\mathbf{x} | U_\theta \mathbf{y}, \mathbf{s})}{p(\mathbf{x} | \mathbf{s})} = I(\mathbf{x}; U_\theta \mathbf{z} | \mathbf{s})$$

(because $p(\mathbf{x}^m | \mathbf{s})$ is independent of m by design of \mathcal{C}). For the GML decoder, we use

$$\begin{aligned} \varphi_{\text{GML}}(\mathbf{x}, \mathbf{y}, \mathbf{s}) &= \max_{\theta' \in \Theta_n} \frac{1}{n} \ln \frac{p(\mathbf{x} | U_{\theta'} \mathbf{y}, \mathbf{s})}{p(\mathbf{x} | \mathbf{s})} \\ &= \max_{\theta' \in \Theta_n} I(\mathbf{x}; U_{\theta'} \mathbf{y} | \mathbf{s}) \end{aligned} \quad (13)$$

The two probabilities in the right side of (12) are given by

$$\begin{aligned} Pr_{\mathbf{X}'}[\varphi_{\text{ML}}(\mathbf{x}', T_\theta \mathbf{z}, \mathbf{s}) \geq \varphi_{\text{ML}}(\mathbf{x}, T_\theta \mathbf{z}, \mathbf{s})] &= Pr[I(\mathbf{x}' | \mathbf{z} | \mathbf{s}) \geq I(\mathbf{x} | \mathbf{z} | \mathbf{s})] \\ &= \sum_{T_{\mathbf{x}' | \mathbf{z} \mathbf{s}} : I(\mathbf{x}' | \mathbf{z} | \mathbf{s}) \geq I(\mathbf{x} | \mathbf{z} | \mathbf{s})} \frac{|T_{\mathbf{x}' | \mathbf{z} \mathbf{s}}|}{|T_{\mathbf{x}' | \mathbf{s}}|} \\ &\geq (n+1)^{-|\mathcal{X}| |\mathcal{Z}| |\mathcal{S}|} 2^{-nI(\mathbf{x}; \mathbf{z} | \mathbf{s})} \end{aligned} \quad (14)$$

and (with $\mathbf{y} = T_\theta \mathbf{z}$)

$$\begin{aligned} Pr[\varphi_{\text{GML}}(\mathbf{x}', T_\theta \mathbf{z}, \mathbf{s}) \geq \varphi_{\text{GML}}(\mathbf{x}, T_\theta \mathbf{z}, \mathbf{s})] &= Pr[\max_{\theta'} I(\mathbf{x}' | U_{\theta'} \mathbf{y} | \mathbf{s}) \geq \max_{\theta'} I(\mathbf{x} | U_{\theta'} \mathbf{y} | \mathbf{s})] \\ &\leq Pr[\max_{\theta'} I(\mathbf{x}' | U_{\theta'} \mathbf{y} | \mathbf{s}) \geq I(\mathbf{x} | U_\theta \mathbf{y} | \mathbf{s})] \\ &\leq \sum_{\theta'} Pr[I(\mathbf{x}' | U_{\theta'} \mathbf{y} | \mathbf{s}) \geq I(\mathbf{x} | U_\theta \mathbf{y} | \mathbf{s})] \\ &\leq \sum_{\theta'} \sum_{T_{\mathbf{x}' | U_{\theta'} \mathbf{y}, \mathbf{s}} : I(\mathbf{x}' | U_{\theta'} \mathbf{y} | \mathbf{s}) \geq I(\mathbf{x} | U_\theta \mathbf{y} | \mathbf{s})} \frac{|T_{\mathbf{x}' | U_{\theta'} \mathbf{y}, \mathbf{s}}|}{|T_{\mathbf{x}' | \mathbf{s}}|} \\ &\leq |\Theta_n| (n+1)^{|\mathcal{S}| |\mathcal{X}| (1+|\mathcal{Z}|)} 2^{-nI(\mathbf{x}; \mathbf{z} | \mathbf{s})}. \end{aligned} \quad (15)$$

Combining (14) and (15) we obtain

$$\begin{aligned} \frac{Pr_{\mathbf{X}'}[\varphi_{\text{GML}}(\mathbf{x}', T_\theta \mathbf{z}, \mathbf{s}) \geq \varphi_{\text{GML}}(\mathbf{x}, T_\theta \mathbf{z}, \mathbf{s})]}{Pr_{\mathbf{X}'}[\varphi_{\text{ML}}(\mathbf{x}', T_\theta \mathbf{z}, \mathbf{s}) \geq \varphi_{\text{ML}}(\mathbf{x}, T_\theta \mathbf{z}, \mathbf{s})]} &\leq |\Theta_n| (n+1)^{|\mathcal{S}| |\mathcal{X}| (1+|\mathcal{Z}|)}. \end{aligned} \quad (16)$$

Applying successively the inequalities (9), (12) and (16), we obtain the desired result, (6).

Theorem 6.3 below extends Theorem 6.1 to the case where Θ_n does not satisfy assumption (A2) but the weaker assumption (A2'). In this case the GML decoder is defined as in (5), except that the maximization over θ is performed over the subset $\tilde{\Theta}_n$, which has subexponential cardinality.

Lemma 6.2: Assume the class of mappings $\{T_\theta, \theta \in \Theta_n\}$ satisfies the condition (1). Then

$$\begin{aligned} \max_{\theta \in \Theta_n} \max_{\mathbf{s}} \max_{\mathbf{x} \in T_{\mathbf{x}^* | \mathbf{s}}} \max_{\mathbf{z}} \min_{\theta' \in \tilde{\Theta}_n} |I(\mathbf{x}; \mathbf{z} | \mathbf{s}) - I(\mathbf{x}; U_{\theta'} T_\theta \mathbf{z} | \mathbf{s})| \\ \leq f(\epsilon_n) \triangleq 4|\mathcal{S}| |\mathcal{X}| |\mathcal{Z}| \epsilon_n \log \epsilon_n^{-1} \end{aligned} \quad (17)$$

which converges to 0 as $n \rightarrow \infty$.

Theorem 6.3: Assume the sequence of index sets Θ_n satisfies Assumption (A2'). For any $\theta \in \Theta_n$, we have

$$\frac{P_e(\theta, \psi_{\text{GML}})}{P_e^*(\theta)} \leq |\tilde{\Theta}_n| (n+1)^{|\mathcal{S}| |\mathcal{X}| (1+2|\mathcal{Z}|)} 2^{nf(\epsilon_n)} \quad (18)$$

and therefore the GML decoder is universal.

Sketch of the Proof. The derivations parallel those of Theorem 6.1. The crucial step is the derivation of the counterpart of (15). By application of Lemma 6.2, we show that (15) is replaced with

$$\begin{aligned} Pr[\varphi_{\text{GML}}(\mathbf{x}', T_\theta \mathbf{z}, \mathbf{s}) \geq \varphi_{\text{GML}}(\mathbf{x}, T_\theta \mathbf{z}, \mathbf{s})] \\ \leq |\Theta_n| (n+1)^{|\mathcal{S}| |\mathcal{X}| (1+2|\mathcal{Z}|)} 2^{-n[-f(\epsilon_n) + I(\mathbf{x}; \mathbf{z} | \mathbf{s})]}. \end{aligned} \quad (19)$$

This leads to (18), where the right side is subexponential in n because $f(\epsilon_n)$ converges to zero.

VII. PUBLIC WATERMARKING

Public watermarking may be modeled as an information transmission problem with side information \mathbf{S} at transmitter only. Here $\mathcal{K} = \emptyset$. Random binning codes are used, whose codewords are sequences in an auxiliary space \mathcal{U}^n . For finite alphabets we use conditionally constant-composition binning codes (conditioned in the type of \mathbf{s}) [10]. Let Λ be the set of all possible types of \mathbf{s} ; we have $|\Lambda| \leq (n+1)^{|\mathcal{S}|}$. To each type $\lambda \in \Lambda$ corresponds an array of codewords with $L(\lambda)$ rows and M columns:

$$\mathcal{C}(\lambda) = \{\mathbf{u}^{lm\lambda}, 1 \leq l \leq L(\lambda), 1 \leq m \leq M\} \quad (20)$$

and two conditional types $\mu_{U|S}$ and $\mu_{X|US}$ defined over the alphabets $\mathcal{U} \times \mathcal{S}$ and $\mathcal{X} \times \mathcal{U} \times \mathcal{S}$, respectively. Also define $\rho(\lambda) = I(\mathbf{u}; \mathbf{s})$, where (\mathbf{u}, \mathbf{s}) is an arbitrary pair of sequences with joint type $\mu_{U|S} \circ \lambda$. The codebook is defined as $\mathcal{C} = \bigcup_{\lambda \in \Lambda} \mathcal{C}(\lambda)$.

Upon observing m and \mathbf{s} , the encoder evaluates the type $\lambda = p_{\mathbf{s}}$ and seeks l such that

$$\mathbf{u}^{lm|\lambda} \in \mathcal{C}(\lambda) \cap \mu_{U|S}(\mathbf{s}).$$

Next, the transmitted sequence \mathbf{x} is drawn from the uniform distribution on the conditional type class $\mu_{X|US}(\mathbf{u}^{lm|\lambda}, \mathbf{s})$.

The ML decoder is the maximizer (over $\lambda \in \Lambda$ and $\mathbf{u} \in \mathcal{C}(\lambda)$) of

$$\varphi_{\text{ML}}(\mathbf{u}, \mathbf{y}) = I(\mathbf{u}; U_{\theta} \mathbf{y}) - \rho(\lambda). \quad (21)$$

The GML decoder is defined as the maximizer (over $\lambda \in \Lambda$ and $\mathbf{u} \in \mathcal{C}(\lambda)$) of

$$\varphi_{\text{GML}}(\mathbf{u}, \mathbf{y}) = \max_{\theta' \in \Theta_n} I(\mathbf{u}; U_{\theta'} \mathbf{y}) - \rho(\lambda). \quad (22)$$

An error is declared if the maximum is not unique.

Theorem 7.1: Under Assumption **(A2)**, we have

$$\max_{\theta \in \Theta_n} \frac{P_e(\theta, \psi_{\text{GML}})}{P_e^*(\theta)} \leq |\Theta_n|(n+1)^{|\mathcal{U}|(1+2|\mathcal{Z}|)}, \quad (23)$$

i.e., the GML decoder is universal.

Remark: the bound (23) is independent of the choice of $L(\lambda)$ and $\rho(\lambda)$.

Sketch of the proof.

Given θ and λ , the conditional probability of error of a decoder ψ is the probability that the transmitted codeword $\mathbf{u} \in \mathcal{C}(\lambda)$ is confused with some other codeword $\mathbf{u}' \in \mathcal{C}(\lambda')$ where $\lambda' \in \Lambda$ and the array elements \mathbf{u} and \mathbf{u}' have different column indices. We have

$$P_e(\theta, \psi) = \mathbb{E}_{\lambda \mathbf{U} \mathbf{Z}} \Pr[\text{error} \mid \lambda, \mathbf{u}, T_{\theta} \mathbf{z}, \psi] \quad (24)$$

and

$$\frac{P_e(\theta, \psi_{\text{GML}})}{P_e^*(\theta)} \leq \max_{\lambda \mathbf{u} \mathbf{z}} \frac{\Pr[\text{error} \mid \lambda, \mathbf{u}, T_{\theta} \mathbf{z}, \psi_{\text{GML}}]}{\Pr[\text{error} \mid \lambda, \mathbf{u}, T_{\theta} \mathbf{z}, \psi_{\text{ML}}]} \quad (25)$$

Both the ML and GML decoding rules may be expressed as

$$\psi(\mathbf{u}, \mathbf{y}) = \operatorname{argmax}_{1 \leq m \leq M} \max_{l, \lambda} \varphi_{\psi}(\mathbf{u}^{lm|\lambda}, \mathbf{y})$$

where φ_{ψ} represents a score function. Then

$$\Pr[\text{error} \mid \lambda, \mathbf{u}, T_{\theta} \mathbf{z}, \psi] = 1 - \prod_{\lambda' \in \Lambda} (1 - \Pr_{\mathbf{U}'|\lambda'}[\varphi_{\psi}(\mathbf{u}', T_{\theta} \mathbf{z}) \geq \varphi_{\psi}(\mathbf{u}, T_{\theta} \mathbf{z})])^{(M-1)L(\lambda')}.$$

Moreover, the ratio of the two probabilities in (25) is upper bounded by

$$\max \left(1, \max_{\lambda' \in \Lambda} \frac{\Pr_{\mathbf{U}'|\lambda'}[\varphi_{\text{GML}}(\mathbf{u}', T_{\theta} \mathbf{z}) \geq \varphi_{\text{GML}}(\mathbf{u}, T_{\theta} \mathbf{z})]}{\Pr_{\mathbf{U}'|\lambda'}[\varphi_{\text{ML}}(\mathbf{u}', T_{\theta} \mathbf{z}) \geq \varphi_{\text{ML}}(\mathbf{u}, T_{\theta} \mathbf{z})]} \right).$$

After some calculations, we obtain the inequalities

$$\begin{aligned} \Pr_{\mathbf{U}'|\lambda'}[\varphi_{\text{ML}}(\mathbf{u}', T_{\theta} \mathbf{z}) \geq \varphi_{\text{ML}}(\mathbf{u}, T_{\theta} \mathbf{z})] \\ \geq (n+1)^{-|\mathcal{U}||\mathcal{Z}|} 2^{-n[I(\mathbf{u}; \mathbf{z}) + \rho(\lambda') - \rho(\lambda)]} \end{aligned}$$

and

$$\begin{aligned} \Pr_{\mathbf{U}'|\lambda'}[\varphi_{\text{GML}}(\mathbf{u}', T_{\theta} \mathbf{z}) \geq \varphi_{\text{GML}}(\mathbf{u}, T_{\theta} \mathbf{z})] \\ \leq |\Theta_n|(n+1)^{|\mathcal{U}|(1+|\mathcal{Z}|)} 2^{-n[I(\mathbf{u}; \mathbf{z}) + \rho(\lambda') - \rho(\lambda)]}. \end{aligned}$$

Combining the inequalities above proves (23).

VIII. CONCLUSION

We have studied a model for geometric attacks on watermarking systems and discovered that for finite alphabets, the GML decoder is universal.

This work admits several extensions. In one of them, the channel $W(z|x)$ is unknown; and another one, the channel W has arbitrary memory, subject to almost-sure distortion constraints, as in [9], [10]. Another extension is the Gaussian case [6], [11].

REFERENCES

- [1] M. Kutter, "Watermarking Resisting to Translation, Rotation and Scaling," *Proc. SPIE*, Boston, Vol. 3528, pp. 423–431, 1998.
- [2] J. J. K. O'Ruanaidh and T. Pun, "Rotation, Scale and Translation Invariant Spread Spectrum Digital Image Watermarking," *Signal Processing*, Vol. 66, No. 3, pp. 303–317, 1998.
- [3] S. Pereira and T. Pun, "Robust Template Matching for Affine Resistant Image Watermarks," *IEEE Trans. on Image Processing*, Vol. 9, No. 6, pp. 1123–1129, June 2000.
- [4] C.-Y. Lin, M. Wu, J. A. Bloom, I. J. Cox, M. L. Miller and Y. M. Lui, "Rotation, Scale, and Translation Resilient Watermarking for Images," *IEEE Trans. on Image Proc.*, Vol. 10, No. 5, pp. 767–782, May 2001.
- [5] M. Feder and A. Lapidoth, "Universal Decoding for Channels with Memory," *IEEE Trans. Information Theory*, Vol. 44, No. 5, pp. 1726–1745, Sep. 1998.
- [6] P. Moulin, "Universal Decoding of Watermarks Under Geometric Attacks," *preprint*, May 2006.
- [7] P. Moulin and J. A. O'Sullivan, "Information-Theoretic Analysis of Information Hiding," *IEEE Trans. on Information Theory*, Vol. 49, No. 3, pp. 563–593, 2003.
- [8] M. Feder and N. Merhav, "Universal Composite Hypothesis Testing: A Competitive Minimax Approach," *IEEE Trans. Information Theory*, Vol. 48, No. 6, pp. 1504–1517, June 2002.
- [9] A. Somekh-Baruch and N. Merhav, "On the Error Exponent and Capacity Games of Private Watermarking Systems," *IEEE Trans. Information Theory*, Vol. 49, No.3, pp. 537–562, March 2003.
- [10] P. Moulin and Y. Wang, "Error Exponents for Channel Coding with Side Information," *Proc. IEEE Information Theory Workshop*, San Antonio, TX, Oct. 2004.
- [11] P. Moulin, "On the Optimal Structure of Watermark Decoders Under Desynchronization Attacks," to appear in *Proc. IEEE Int. Conf. on Image Proc.*, Atlanta, GA, Oct. 2006.